



US009491892B1

(12) **United States Patent**
Carlson et al.

(10) **Patent No.:** **US 9,491,892 B1**
(45) **Date of Patent:** **Nov. 8, 2016**

(54) **ELECTRONIC DEVICE COOLING SYSTEM WITH STORAGE**

(71) Applicant: **Google Inc.**, Mountain View, CA (US)

(72) Inventors: **Andrew B. Carlson**, Atherton, CA (US); **William Hamburg**, Palo Alto, CA (US); **Donald L. Beaty**, Red Bank, NJ (US); **Robert Wasilewski**, Hightstown, NJ (US)

(73) Assignee: **Google Inc.**, Mountain View, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 328 days.

(21) Appl. No.: **14/047,574**

(22) Filed: **Oct. 7, 2013**

Related U.S. Application Data

(63) Continuation of application No. 12/060,222, filed on Mar. 31, 2008, now Pat. No. 8,553,416.

(60) Provisional application No. 61/016,419, filed on Dec. 21, 2007.

(51) **Int. Cl.**
H05K 7/20 (2006.01)
F24J 2/30 (2006.01)
F24J 2/10 (2006.01)

(52) **U.S. Cl.**
CPC **H05K 7/20836** (2013.01); **F24J 2/10** (2013.01); **F24J 2/30** (2013.01); **H05K 7/20381** (2013.01)

(58) **Field of Classification Search**
None
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,850,007 A 11/1974 McFarlan
3,995,443 A 12/1976 Iversen

4,010,624 A	3/1977	McFarlan
4,144,722 A	3/1979	Mattwell
4,285,392 A	8/1981	Rannow
4,313,310 A	2/1982	Kobayashi et al.
4,315,404 A *	2/1982	Schmitt F28C 1/00 165/104.14
4,457,358 A	7/1984	Kriege et al.
4,474,027 A	10/1984	Kaya et al.
4,488,408 A	12/1984	Kajitsuka
4,931,187 A	6/1990	Derham et al.
5,057,229 A	10/1991	Schulenburg et al.
5,145,585 A *	9/1992	Coke C02F 1/385 210/167.29
5,730,879 A	3/1998	Wilding et al.
5,858,219 A	1/1999	Kusmierz et al.
5,963,458 A	10/1999	Cascia
5,970,729 A	10/1999	Yamamoto et al.
6,012,708 A	1/2000	Nagano
6,142,219 A	11/2000	Korenic et al.
6,213,200 B1	4/2001	Carter et al.
6,418,728 B1	7/2002	Monroe
6,446,448 B1	9/2002	Wang et al.
6,446,941 B1	9/2002	Maheshwari et al.
6,532,754 B2	3/2003	Haley et al.
6,564,864 B2	5/2003	Carter et al.
6,604,376 B1	8/2003	Demarco et al.
6,967,842 B2	11/2005	Aoki et al.

(Continued)

FOREIGN PATENT DOCUMENTS

JP 6313646 A 11/1994
JP 2002262839 A 9/2002

(Continued)

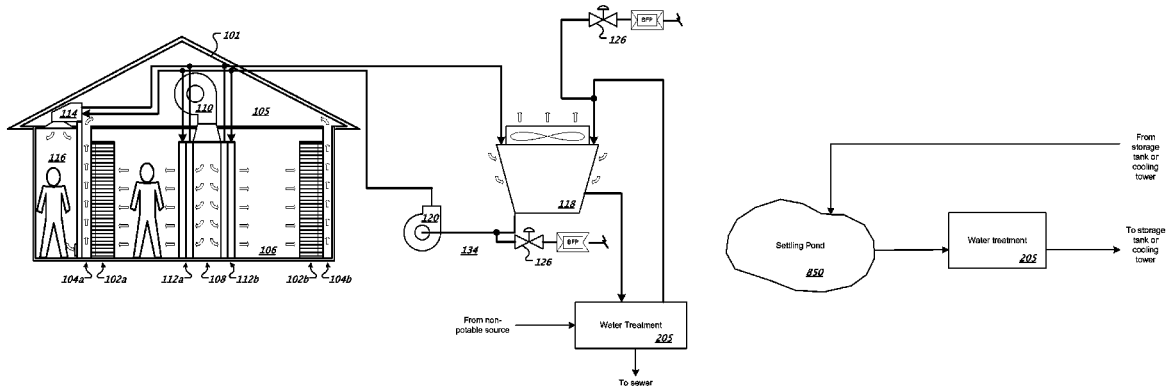
Primary Examiner — Courtney Smith

(74) *Attorney, Agent, or Firm* — Fish & Richardson P.C.

(57) **ABSTRACT**

Cooling systems for providing cooled air to electronic equipment are described. The systems can include large storage tanks or waste treatment systems to improve the efficiency of the plant and reduce impact on the environment.

7 Claims, 10 Drawing Sheets



(56)

References Cited

U.S. PATENT DOCUMENTS

6,973,801	B1	12/2005	Campbell et al.	
7,184,269	B2	2/2007	Campbell et al.	
7,195,176	B2	3/2007	Newman	
7,215,545	B1	5/2007	Moghaddam et al.	
7,340,912	B1	3/2008	Yoho et al.	
7,511,959	B2	3/2009	Belady et al.	
7,511,960	B2	3/2009	Hillis et al.	
7,660,117	B2	2/2010	Werner et al.	
7,724,513	B2	5/2010	Coglitore et al.	
7,746,634	B2	6/2010	Hom et al.	
7,788,941	B2	9/2010	Campbell et al.	
7,854,652	B2	12/2010	Yates et al.	
7,856,838	B2	12/2010	Hillis et al.	
7,864,530	B1	1/2011	Hamburgien et al.	
8,223,495	B1 *	7/2012	Carlson	F24F 5/0035 165/104.33
8,635,881	B2 *	1/2014	Carlson	F24F 11/0001 62/168
8,667,804	B2 *	3/2014	Koh	F24F 5/0017 165/10
2002/0014460	A1	2/2002	McKay	
2002/0017095	A1	2/2002	Pierson	
2002/0053214	A1	5/2002	Melendez-Gonzalez et al.	
2003/0029200	A1	2/2003	Mace	
2003/0029390	A1	2/2003	Campion	
2003/0037905	A1	2/2003	Weng	
2003/0173688	A1	9/2003	Koo	
2004/0060808	A1 *	4/2004	LaViolette	B01D 3/346 202/234
2004/0089011	A1	5/2004	Patel et al.	
2004/0206101	A1	10/2004	Bash et al.	
2005/0056023	A1	3/2005	Pierson	
2005/0072171	A1	4/2005	Jensen	
2005/0103032	A1	5/2005	Pierson	
2005/0241799	A1	11/2005	Malone et al.	
2005/0241803	A1	11/2005	Malone et al.	
2005/0247638	A1	11/2005	Owens	
2006/0082263	A1	4/2006	Rimler et al.	
2006/0260605	A1 *	11/2006	Connor	F24J 2/085 126/561
2007/0074525	A1	4/2007	Vinson et al.	
2007/0076373	A1	4/2007	Fink	
2007/0102367	A1	5/2007	Owens	
2007/0151278	A1	7/2007	Jarvis	
2007/0163748	A1	7/2007	Rasmussen et al.	
2007/0165377	A1	7/2007	Rasmussen et al.	

2007/0167125	A1	7/2007	Rasmussen et al.	
2007/0177351	A1	8/2007	Monfarad et al.	
2008/0023409	A1	1/2008	Owens	
2008/0029250	A1*	2/2008	Carlson	F24F 11/0001 165/104.33
2008/0055846	A1	3/2008	Clidasar et al.	
2008/0060372	A1	3/2008	Hillis et al.	
2008/0062647	A1	3/2008	Hillis et al.	
2008/0064317	A1	3/2008	Yates et al.	
2008/0094797	A1	4/2008	Coglitore et al.	
2008/0123288	A1	5/2008	Hillis	
2008/0198549	A1	8/2008	Rasmussen et al.	
2008/0204999	A1	8/2008	Clidasar et al.	
2008/0209931	A1	9/2008	Stevens	
2008/0229780	A1	9/2008	Wyatt et al.	
2008/0259566	A1	10/2008	Fried	
2008/0266794	A1	10/2008	Malone	
2008/0270572	A1	10/2008	Belady et al.	
2008/0307806	A1	12/2008	Campbell et al.	
2009/0086434	A1	4/2009	Hodes et al.	
2009/0090488	A1	4/2009	McCann	
2009/0097205	A1	4/2009	Matsushima et al.	
2009/0126910	A1	5/2009	Campbell et al.	
2009/0171512	A1	7/2009	Duncan	
2009/0179429	A1	7/2009	Ellis et al.	
2009/0295167	A1	12/2009	Clidasar et al.	
2009/0296346	A1	12/2009	Van Anel	
2010/0018228	A1*	1/2010	Flammang	F25B 27/02 62/115
2010/0041327	A1	2/2010	Desler	
2010/0043473	A1*	2/2010	Koh	F24F 5/0017 62/259.1
2010/0236266	A1*	9/2010	Skidmore	F24J 3/08 62/238.7
2011/0120693	A1	5/2011	Kammerzell et al.	
2011/0137479	A1*	6/2011	Al-Mazeedi	G05B 13/021 700/287
2011/0277967	A1	11/2011	Fried et al.	
2012/0255706	A1*	10/2012	Tadayon	F24J 3/081 165/47
2013/0047653	A1*	2/2013	Keisling	H05K 7/20827 62/228.7

FOREIGN PATENT DOCUMENTS

JP	2003121024	A	4/2003
WO	WO2007139558	A1	12/2007

* cited by examiner

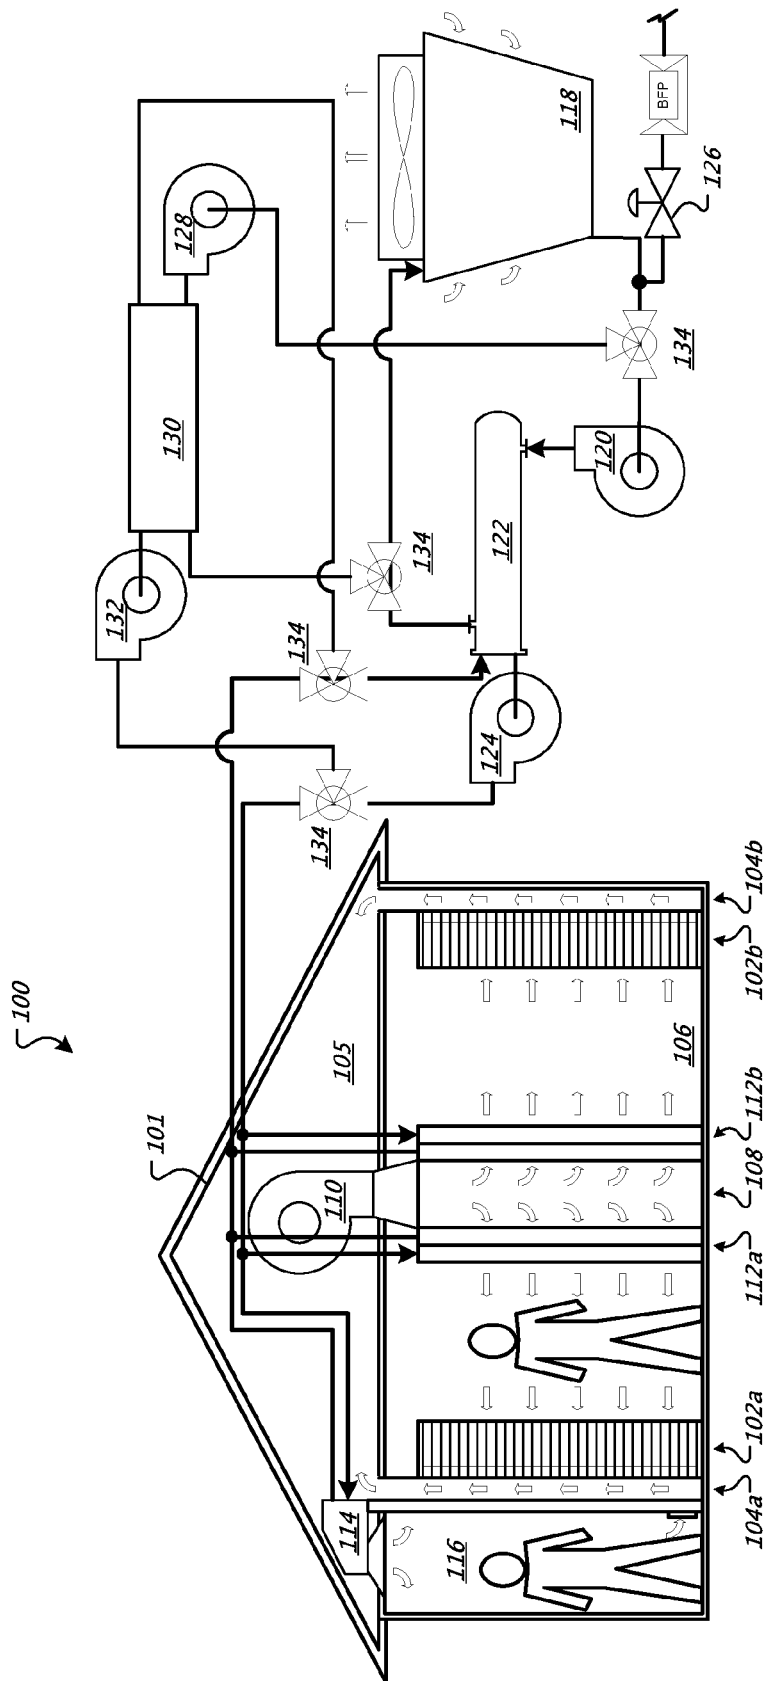


FIG. 1

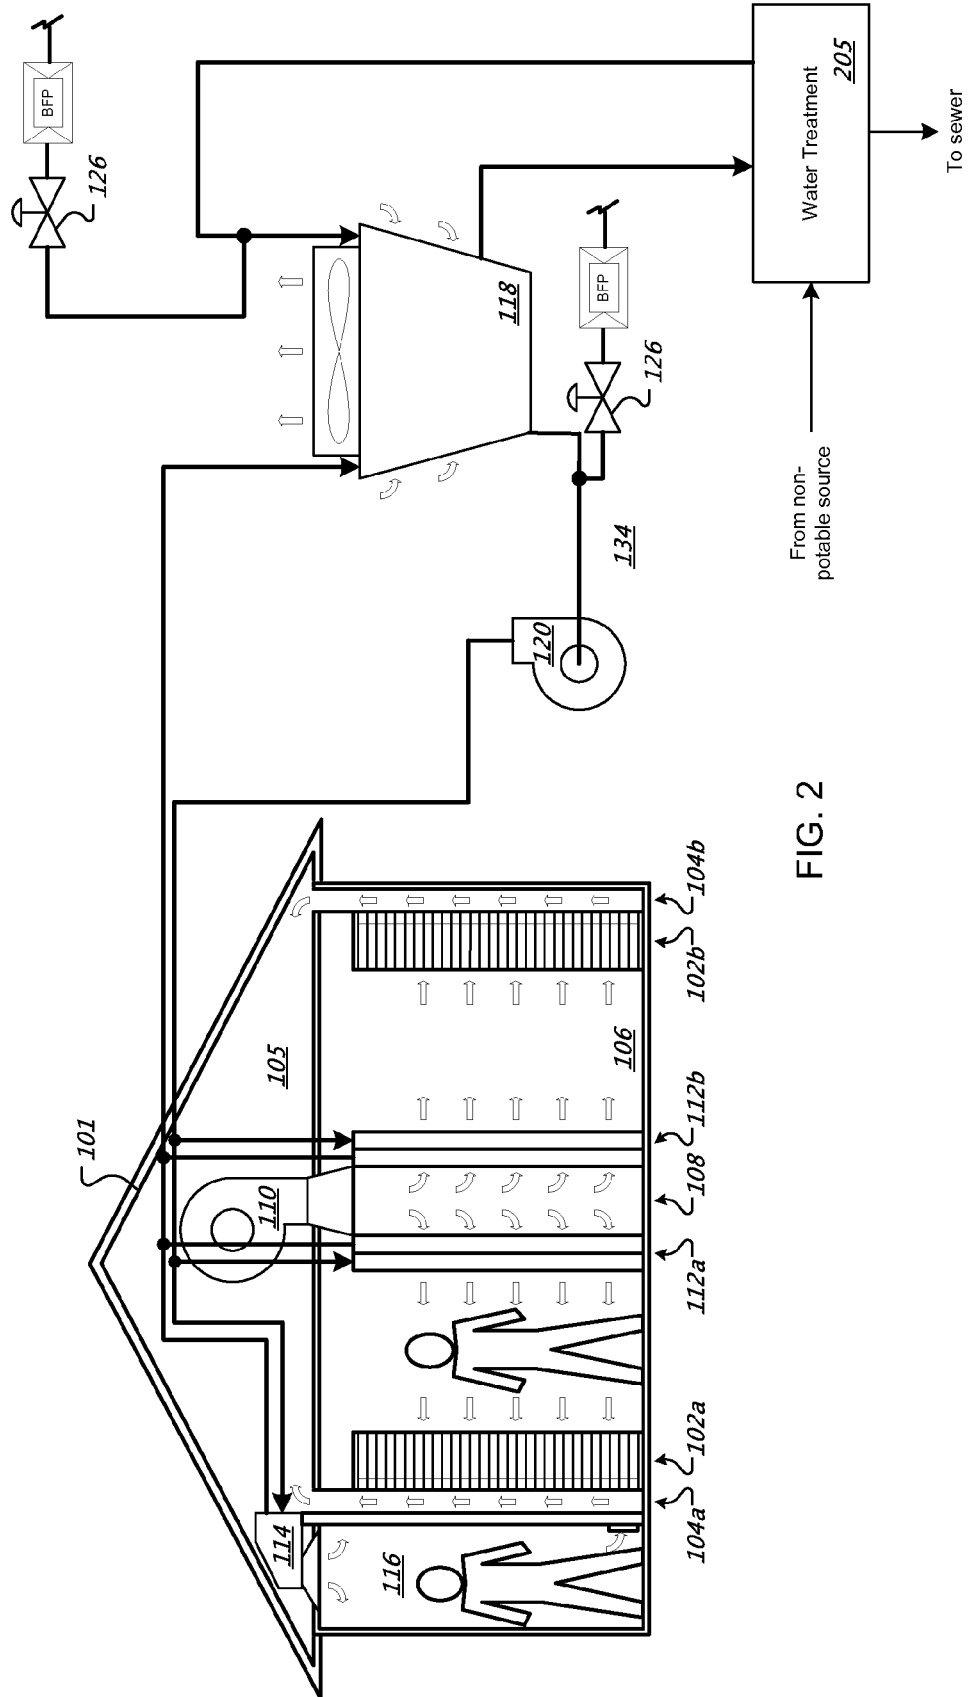
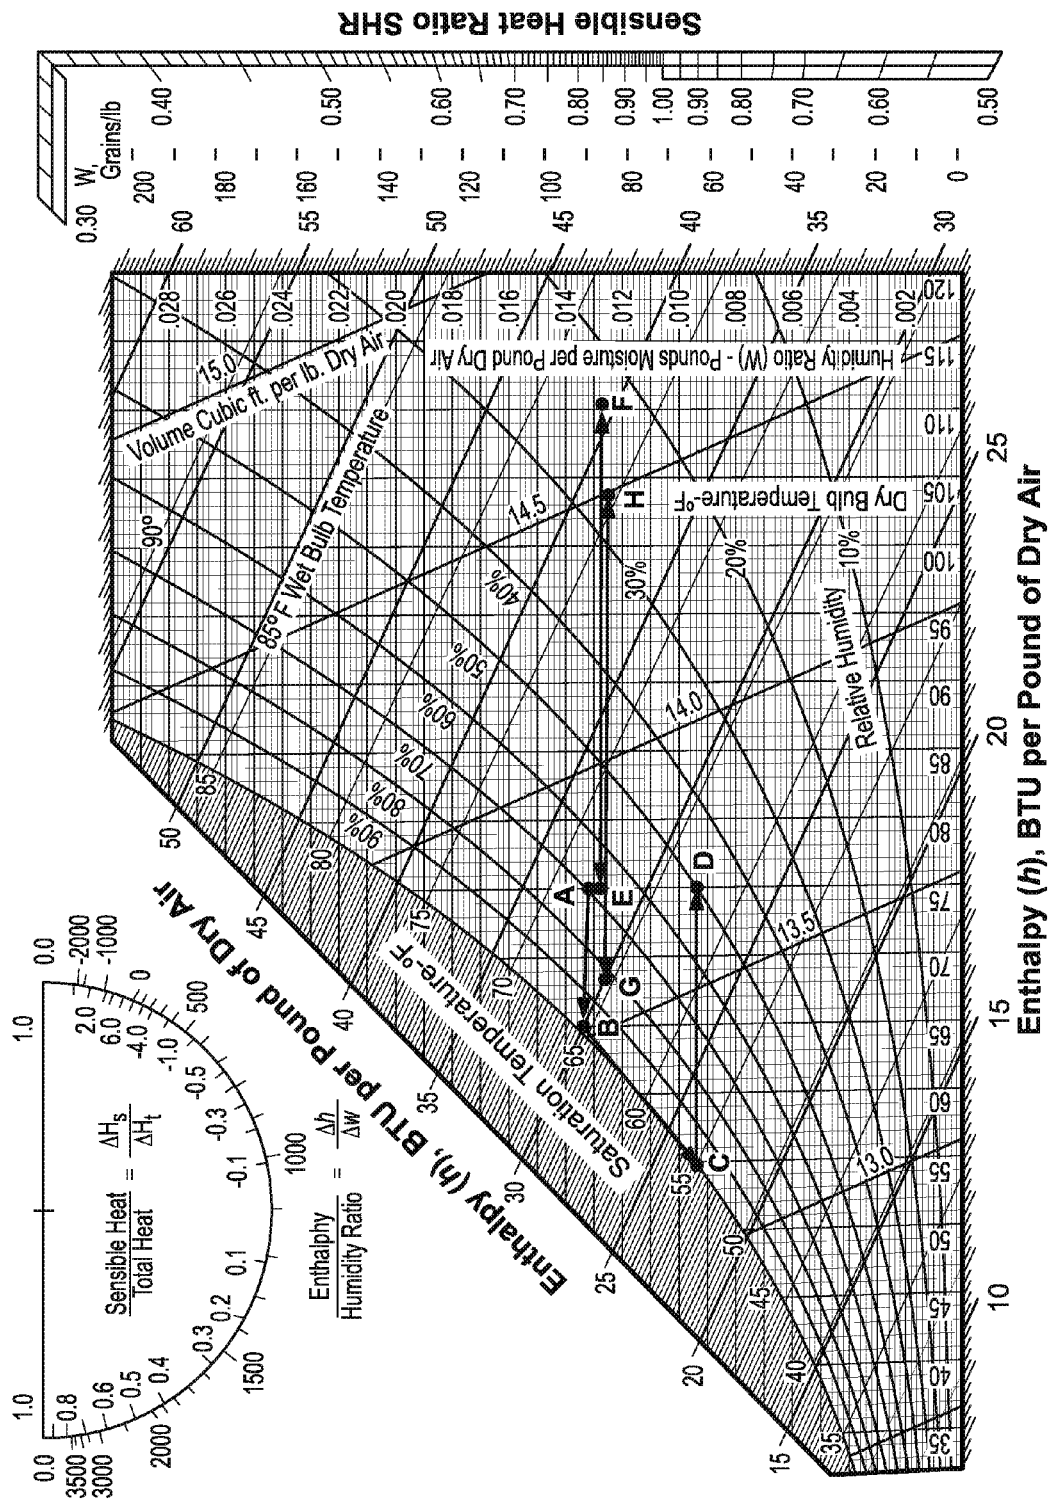


FIG. 2



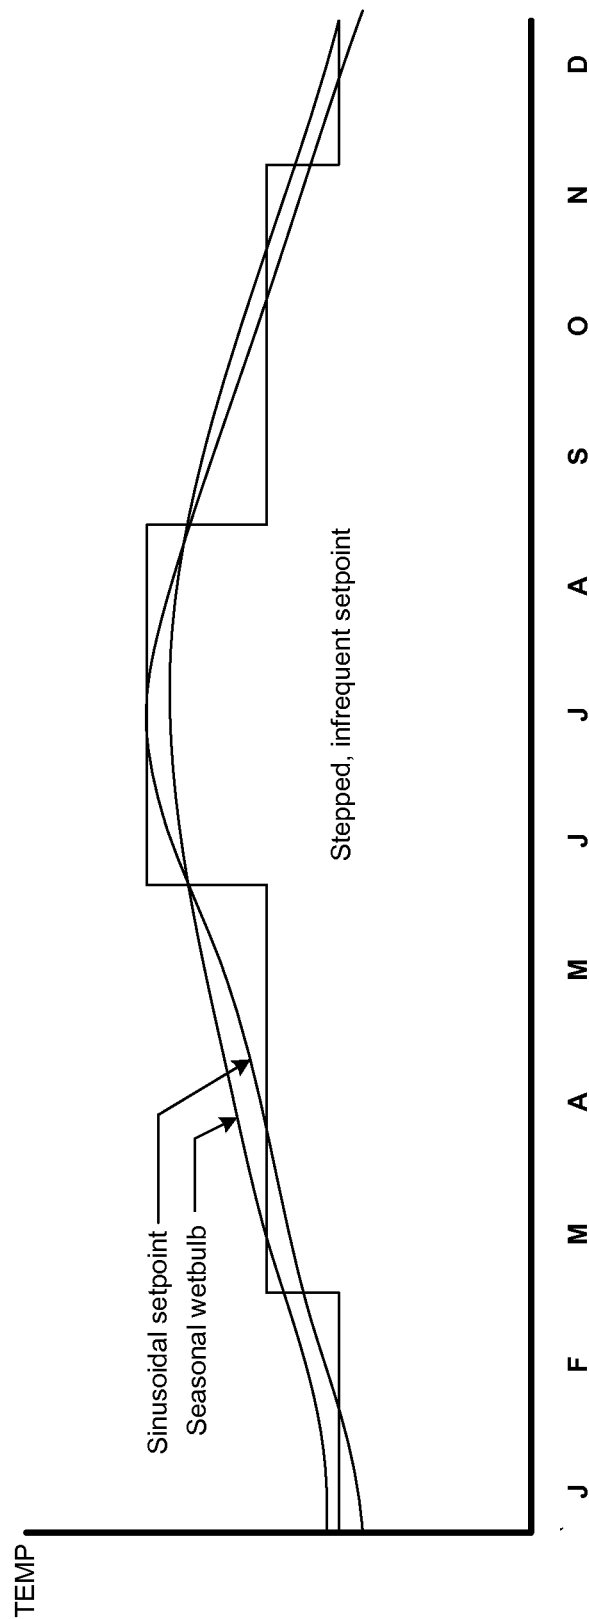


FIG. 3B

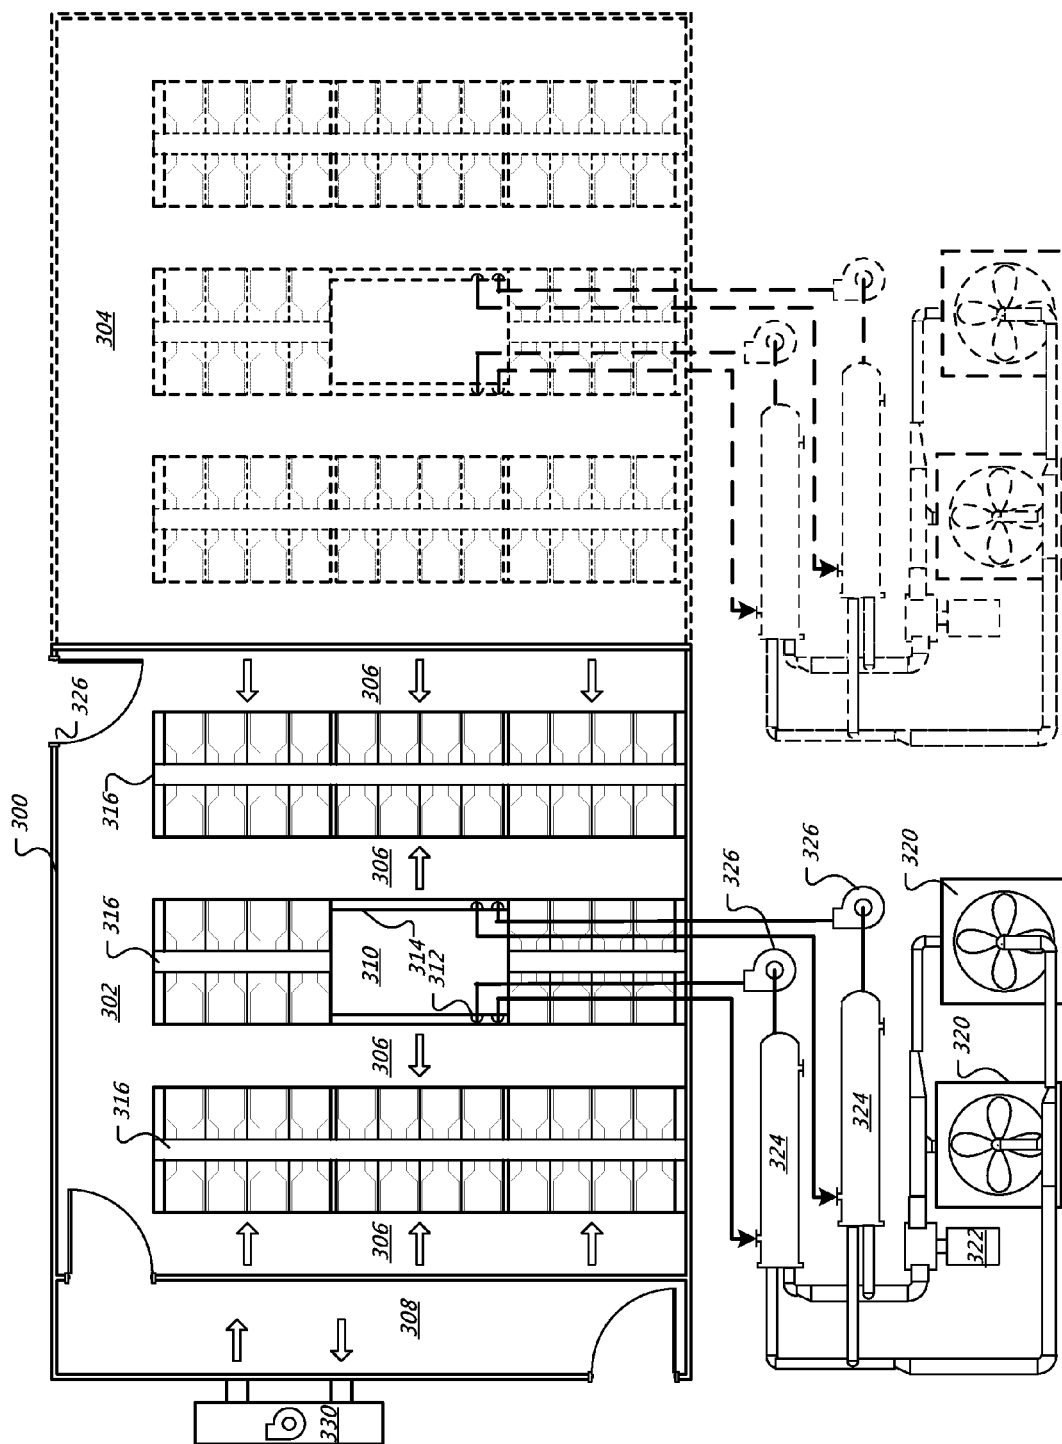


FIG. 4

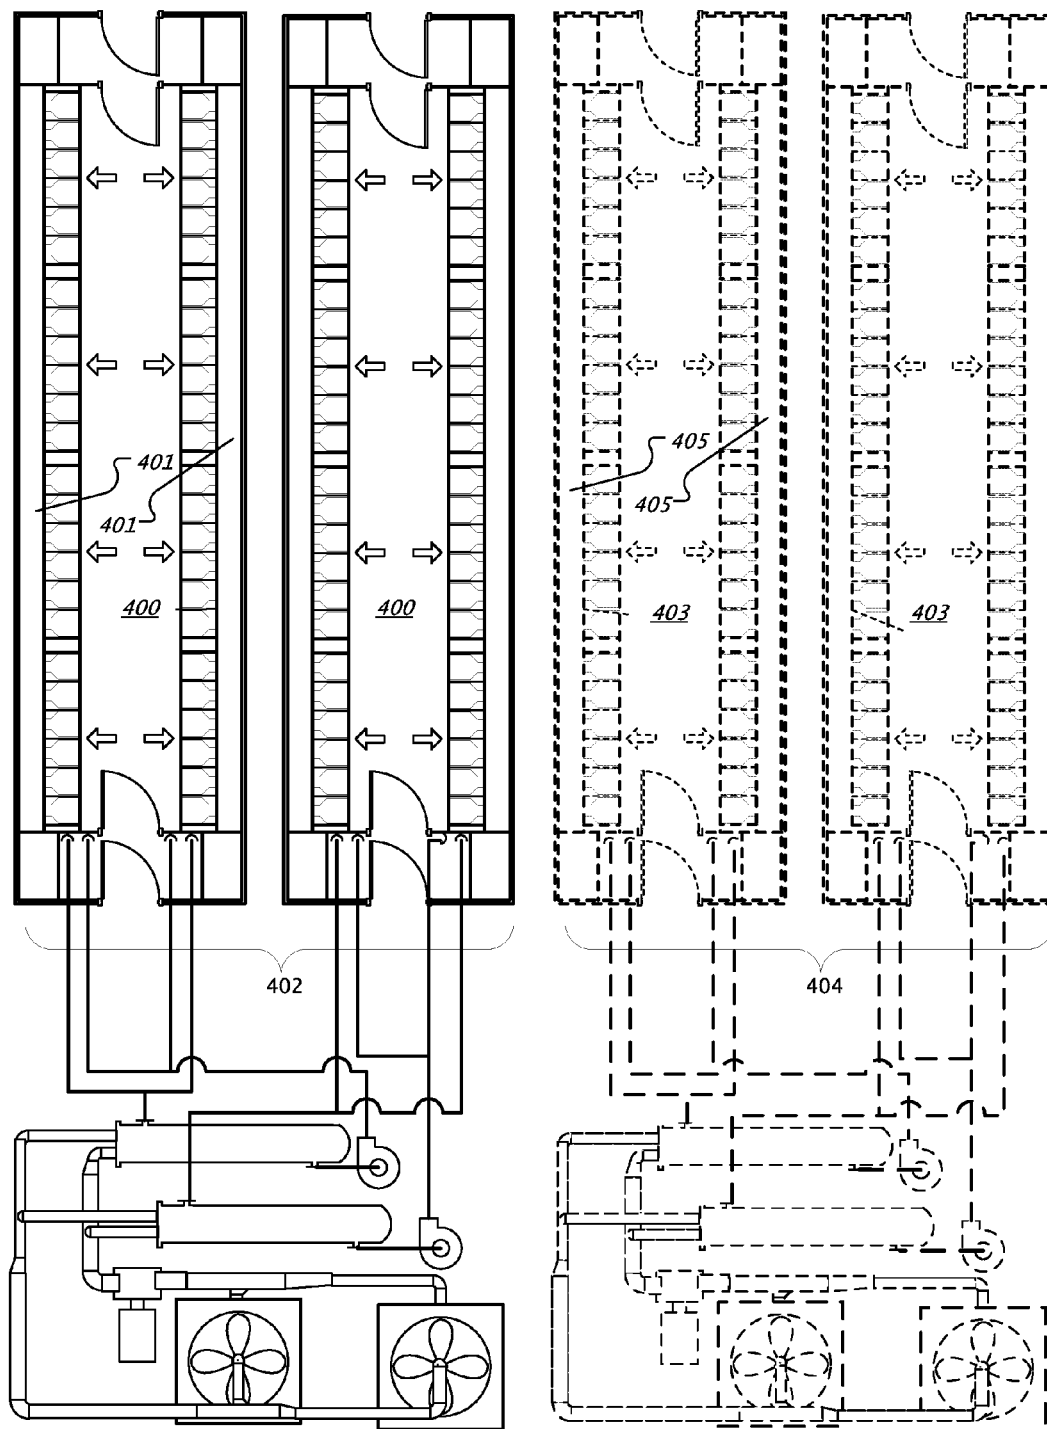


FIG. 5

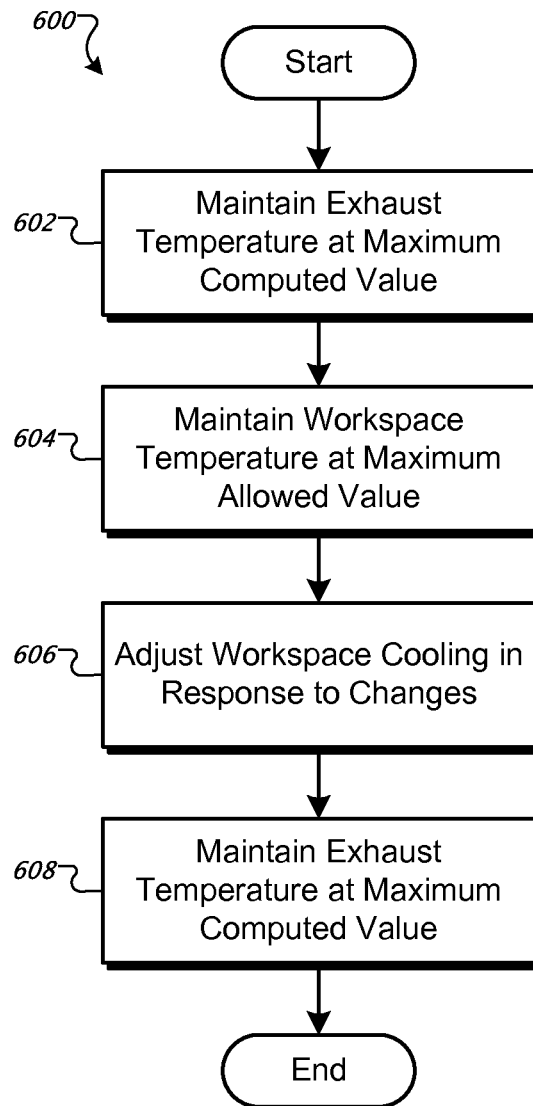


FIG. 6

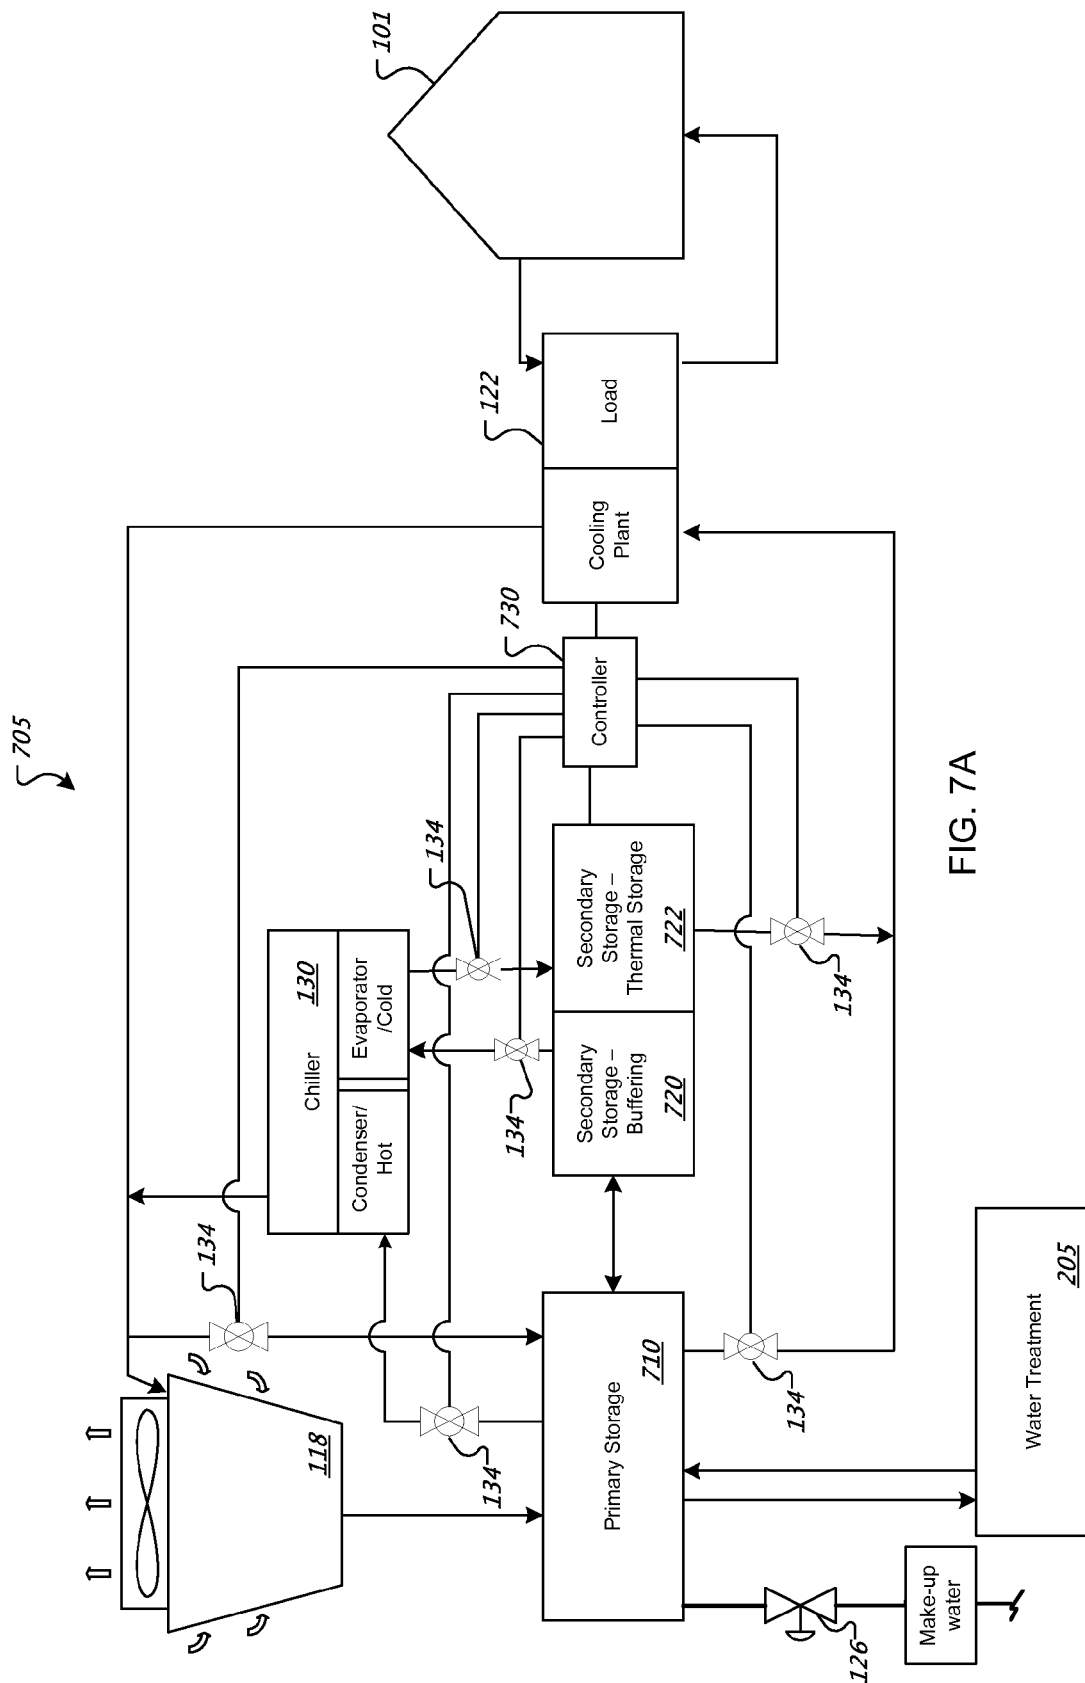


FIG. 7A

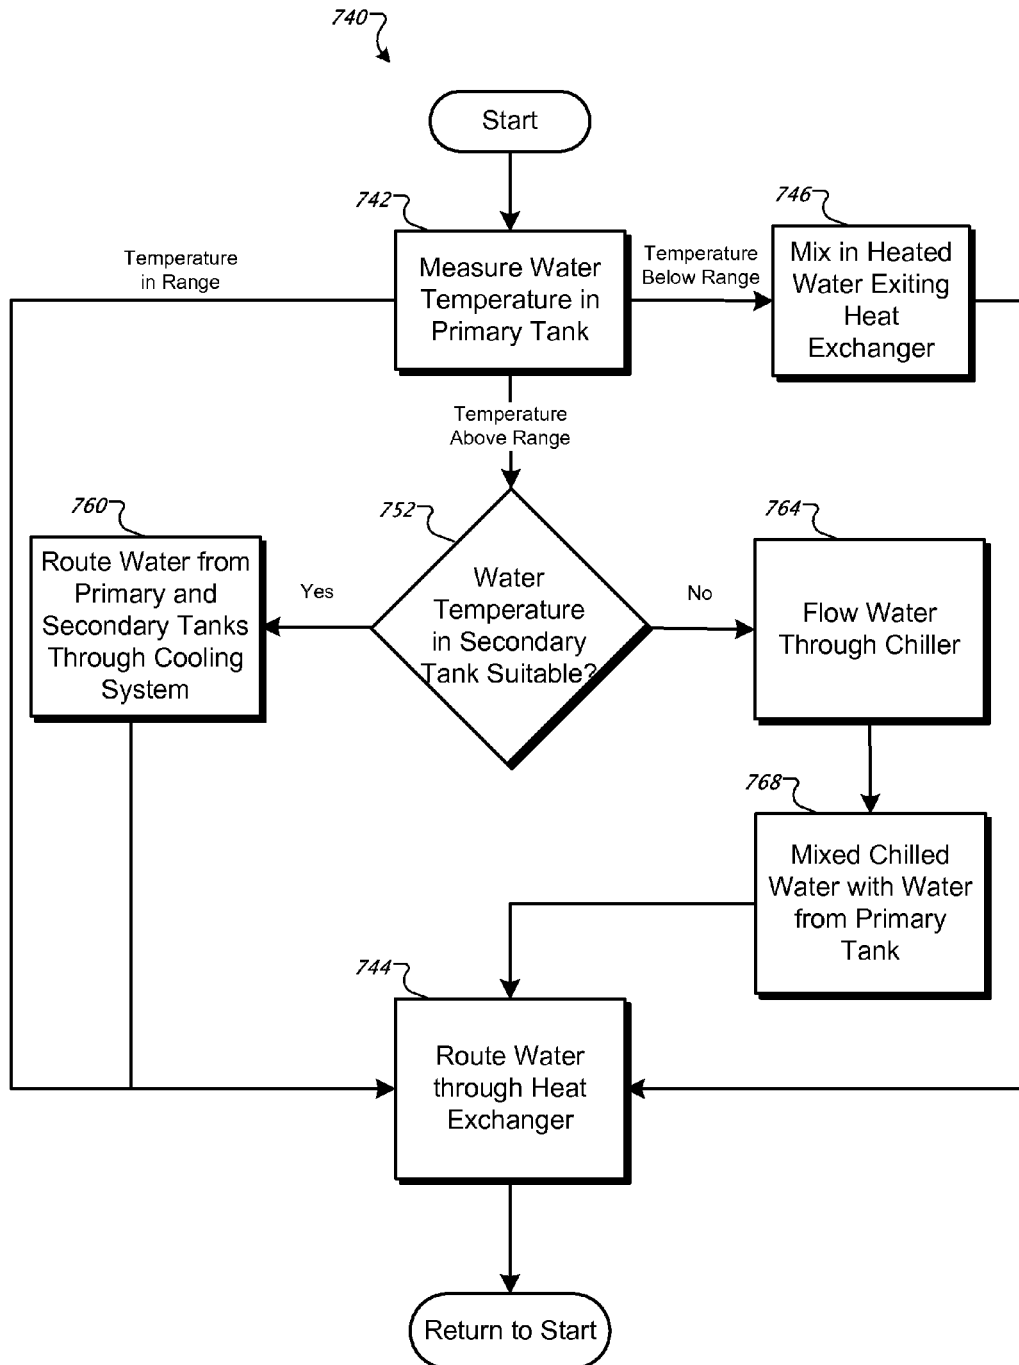


FIG. 7B

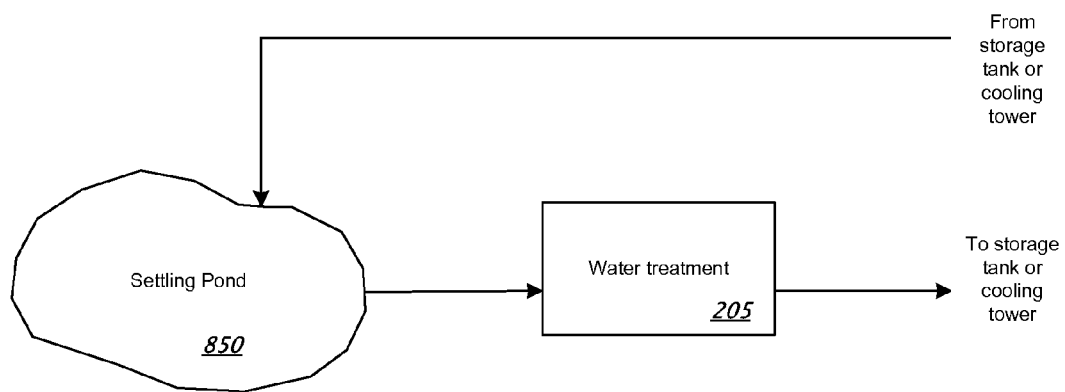


FIG. 8

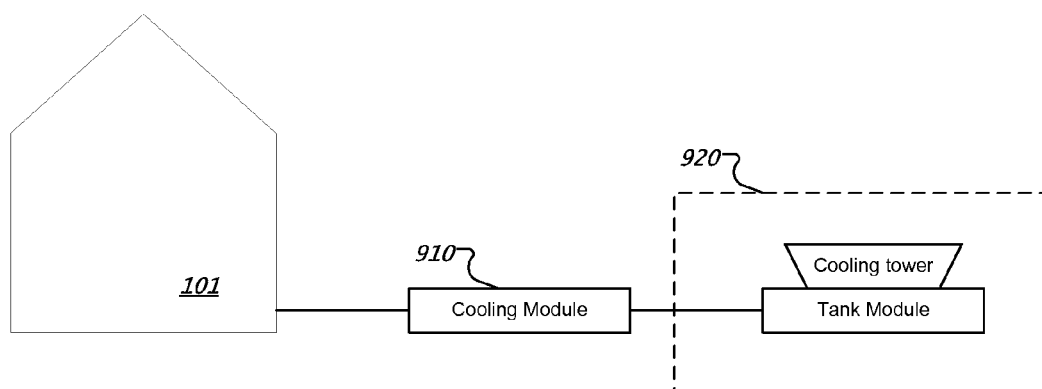


FIG. 9

1

ELECTRONIC DEVICE COOLING SYSTEM WITH STORAGE

CROSS-REFERENCE TO RELATED APPLICATIONS

This application is a continuation application and claims the benefit of priority under 35 U.S.C. Section 120 of U.S. application Ser. No. 12/060,222, filed on Mar. 31, 2008, which claims priority to U.S. Provisional Application Ser. No. 61/016,419, filed on Dec. 21, 2007. The disclosure of each prior application is considered part of and is incorporated by reference in the disclosure of this application.

TECHNICAL FIELD

This document relates to systems and methods for providing cooling for areas containing electronic equipment, such as computer server rooms and server racks in computer data centers.

BACKGROUND

Computer users often focus on the speed of computer microprocessors (e.g., megahertz and gigahertz). Many forget that this speed often comes with a cost—higher electrical power consumption. For one or two home PCs, this extra power may be negligible when compared to the cost of running the many other electrical appliances in a home. But in data center applications, where thousands of microprocessors may be operated, electrical power requirements can be very important.

Power consumption also, in effect, creates a double hit. Not only must a data center operator pay for electricity to operate its many computers, but the operator must also pay to cool the computers. That is because, by simple laws of physics, all the power has to go somewhere, and that somewhere is, for the most part, conversion into heat. A pair of microprocessors mounted on a single motherboard can draw 200-400 watts or more of power. Multiply that figure by several thousand (or tens of thousands) to account for the many computers in a large data center, and one can readily appreciate the amount of heat that can be generated. It is much like having a room filled with thousands of burning floodlights.

Thus, the cost of removing all of the heat can also be a major cost of operating large data centers. That cost typically involves the use of even more energy, in the form of electricity and natural gas, to operate chillers, condensers, pumps, fans, cooling towers, and other related components. Heat removal can also be important because, although microprocessors may not be as sensitive to heat as are people, increases in heat generally can cause great increases in microprocessor errors and failures. In sum, such a system may require electricity to run the chips, and more electricity to cool the chips.

SUMMARY

This document describes systems and methods that may be employed to remove heat efficiently from areas hosting electronic equipment, such as data centers. In certain implementations, cooling of equipment may occur at elevated air temperatures. For example, high temperature rises may be created across a group of heat-generating components by intentionally slowing the flow of cooling air over the components. As one example, temperature rises across the

2

components of thirty-six degrees Fahrenheit (twenty degree Celsius) or more may be maintained, with entering air temperatures of about seventy-seven degrees Fahrenheit (twenty-five degrees Celsius) and exiting temperatures of about one hundred thirteen degrees Fahrenheit (forty-five degrees Celsius). The temperatures of any associated cooling water may also be elevated, such as to produce supply temperatures of about sixty-eight degrees Fahrenheit (twenty degrees Celsius) and return temperatures of about one hundred four degrees Fahrenheit (forty degrees Celsius). Because of the elevated temperatures, the system may be run under most conditions using only cooling from cooling towers or other free cooling sources, without the need for a chiller or other similar sources that require additional power for cooling.

In one embodiment, a system for providing cooled air to electronic equipment is described. The system has a cooling tower, a water storage tank fluidly coupled to the cooling tower and an air-to-water heat exchanger thermally connected to the water storage tank and positioned to receive heated air from a group of electronic devices.

In another embodiment, a method of providing cooled air to electronic equipment is described. Water in a cooling tower is allowed to cool through evaporative cooling. Water is flowed from the cooling tower to a water storage tank. Water is flowed from the water storage tank to an air-to-water heat exchanger. The water flowing through the air-to-water heat exchanger cools air surrounding the heat exchanger that has been heated by a group of electronic devices.

In yet another embodiment, a method of providing cooled air to electronic equipment is described. Water in a cooling tower is allowed to cool through evaporative cooling. Water is flowed from the cooling tower to a water storage tank. Water is flowed from the water storage tank to a water-to-water heat exchanger. Water is flowed from the water-to-water heat exchanger through an air-to-water heat exchanger. The water flowing through water-to-air heat exchanger is kept isolated from the water from the water storage tank and is flowed through the air-to-water heat exchanger to cool air surrounding the heat exchanger that has been heated by a group of electronic devices.

In another embodiment, a system for providing cooled air to electronic equipment is described. The system includes a cooling tower, a water treatment system fluidly coupled to the cooling tower and an air-to-water heat exchanger connected by a pair of pipes to the cooling tower and positioned to receive heated air from a group of electronic devices.

In yet another embodiment, a method of providing cooled air to electronic equipment is described. Water in a cooling tower is allowed to cool through evaporative cooling. Water from the cooling tower is flowed to an air-to-water heat exchanger, wherein the water flowing through the air-to-water heat exchanger cools air surrounding the heat exchanger that has been heated by a group of electronic devices. Water heated by the air surrounding the heat exchanger is flowed into the cooling tower. Water is flowed from the cooling tower into a water treatment system for removing impurities from the water.

In another embodiment, a system for providing cooled air to electronic equipment is described. The system includes a source of non-potable water, a cooling tower in fluid communication with the source of non-potable water, and an air-to-water heat exchanger thermally connected to the cooling tower and positioned to receive heated air from a group of electronic devices.

In yet another embodiment, method of providing cooled air to electronic equipment is described. Non-potable water is provided to a cooling tower. Water in the cooling tower is allowed to cool through evaporative cooling. Water is flowed from the cooling tower to an air-to-water heat exchanger, wherein the water flowing through the air-to-water heat exchanger cools air surrounding the heat exchanger that has been heated by a group of electronic devices.

In another embodiment, a system for providing cooled air to electronic equipment is described. The system includes a water storage pond, a reflective cover on the storage pond, a water treatment system fluidly coupled to the storage pond and an air-to-water heat exchanger connected to the water treatment system and positioned to receive heated air from a group of electronic devices.

In yet another embodiment, a method of providing cooled air to electronic equipment. Water is provided to a water storage pond. Water in the storage pond is allowed to cool through evaporative cooling. Water is flowed from the storage pond to a water treatment system and impurities are removed from the water at the water treatment system. Water is flowed from the water treatment system to an air-to-water heat exchanger, wherein the water flowing through the air-to-water heat exchanger cools air surrounding the heat exchanger that has been heated by a group of electronic devices.

In another embodiment, a modular system for providing cooling water to a building housing a group of electronic devices includes a module and a water storage. The module includes a housing, wherein the housing is no more than 13.5 feet wide and is no longer than 97 feet, a water-to-water heat exchanger within the housing, a chiller within the housing and a controller within the housing, wherein the controller regulates a flow of water from the water-to-water heat exchanger and the chiller to outside of the housing. The water storage tank has a capacity of at least 10,000 gallons and is fluidly coupled to at least one of the chiller and the water-to-water heat exchanger.

Advantages of the systems and methods described herein may include one or more of following. Recirculation of water through a cooling system can decrease the system's reliance on municipal or county water sources. This may both reduce the cost of operation of the cooling system, as well as benefit the local municipality. The municipality's resources are not as heavily drawn upon by the cooling system as with other types of water based cooling systems, such as systems that run tap water into the system and straight in to the sewer after the water has been used. Further, a system as described herein can reduce the load upon the local sewer system and sewage treatment plant, as only a fraction of the water used by the cooling system may be released into the sewage system each day. An optional on-site water treatment facility may further reduce the load on the local sewage system. The on-site water treatment facility may be used to treat water coming into the system that is not as clean as needed and may even provide environmental benefits of cleaning water that would normally not be treated, and later returning the water to the environment cleaner than when it was removed. If the cooling system includes a storage tank, the storage tank can provide a source of water for cooling in the event the regularly available water is out of service for some period of time, such as for up to a day or a few days. Further, the cooling system can provide massive storage that may be able to compensate for any air temperature peaks and troughs that occur over some period of time, such as over a day or over

a few days. The tank system may also reduce the need for chillers to chill the water, which in turn can reduce the amount of electricity required by a cooling plant. The use of less electricity may enable locating the cooling system in a greater variety of locations than a system that has high electrical needs. A cooling system with lower electrical needs may be both more cost efficient to operate and have a lower environmental impact than more energy intensive plants. The cooling system can be modularized. This can allow for construction of complex piping and wiring to be performed in a controlled environment, such as a factory, rather than on-site. Controlled build conditions can increase repeatability and decrease construction problems. The modules also enable a data center to be built at a particular capacity and expanded to greater capacity as the need for more electronic equipment, such as servers, grows. Further, this allows for scaling the cost of building the data center to the current needs of the data center, while leaving room for future growth.

The details of one or more embodiments are set forth in the accompanying drawings and the description below. Other features, objects, and advantages will be apparent from the description and drawings, and from the claims.

DESCRIPTION OF DRAWINGS

FIG. 1 is a schematic diagram showing a system for cooling a computer data center.

FIG. 2 is a schematic diagram showing a system for cooling a computer data center.

FIG. 3A is a psychometric chart showing a heating and cooling cycle for air in a data center.

FIG. 3B is a graph of setpoint temperature for a computing facility over a one year time period.

FIG. 4 is a plan view in schematic form of a hydronic cooling system for a data center.

FIG. 5 is a plan view in schematic form of a hydronic cooling system for a data center housed in shipping containers.

FIG. 6 is a flowchart showing steps for cooling a data center using elevated temperatures.

FIG. 7A is a schematic diagram showing a system for cooling a computer data center.

FIG. 7B is a flowchart showing the steps for routing water through the cooling system.

FIG. 8 is a schematic diagram including a settling pond.

FIG. 9 is a schematic diagram of a data center with one modularized cooling system.

Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

FIG. 1 is a schematic diagram showing a system 100 for cooling a computer data center 101. The system 100 generally includes an air handling unit (including e.g., fan 110 and cooling coils 112a, 112b) in the data center 101 for transferring heat from the data center's air to cooling water, a heat exchanger 122 for removing heat from the cooling water and passing it to cooling tower water, and a cooling tower 118 to pass the accumulated heat to the ambient air through evaporation and cooling of the cooling tower water. In general operation, the system 100 may be run from the cooling tower/heat exchanger/cooling coil system, though a powered refrigeration system such as a chiller may be provided for peak loads, such as when the outdoor ambient dew point is very high and the cooling tower cannot provide

5

sufficient cooling alone. As explained below, control parameters for the system may also be set so as to avoid most or any need for the use of chillers or other such items.

The temperatures of each portion of the system **100** are selected to be relatively high, so as to permit more efficient operation of the system **100**, than if the temperatures were lower. For example, relatively high air temperatures in the system (e.g., air entering a cooling coil over 110° F. (43.3° C.) and exiting temperature above 70° F. (21.11° C.)) may in turn permit for relatively high cooling water temperatures (e.g., water entering a cooling coil around 68° F. (20° C.) and exiting around 104° F. (40° C.)) because the amount of heat that can be taken out of the air is generally proportional to the difference in temperature between the water and the air. If the difference can be kept at an acceptable level, where the temperatures are high enough that evaporative cooling (e.g., cooling through a cooling tower, without further cooling via chiller) is sufficient, the relatively high electrical cost of operating a chiller (or many chillers) may be avoided.

High system temperatures may be particularly advantageous in certain implementations when hybrid cooling towers are used. Such hybrid cooling towers combine the functionality of an ordinary cooling tower with a water-to-water heat exchanger. Sufficiently high chosen temperature setpoints may allow the hybrid tower to provide substantial cooling capacity, even when operating in a water-to-air mode without utility water. As a result, a hybrid cooling tower may be used to provide cooling capacity to a facility relatively quickly, even before utility water may be obtained in large volumes. The capacity of the cooling tower may be directly related to the difference in the temperature of the water within it to the ambient outside air.

When the difference in temperatures is not very large, a change of only a few degrees can bring substantial gains in efficiency. For example, where the cooling water enters at 68° F. (20° C.), by heating air to 113° F. (45° C.) rather than 104° F. (40° C.), the temperature difference is increased from 36° F. to 45° F. (2° C. to 7° C.)—which may result in an increase in heat flow of 25 percent. The actual difference will vary slightly, as the entering conditions for air and water are not the only conditions (because the air cools as it passes through a cooling coil, and the water warms); this example, however, indicates how the difference in temperature can affect efficiency of a system.

Use of elevated temperatures in a system may also prevent air in or around the system from falling below its liquid saturation point, i.e., its dew point, and condensing. This may, in certain circumstances, provide benefits both in efficiency and in operations of the system. Efficiency benefits may be obtained because creating condensation requires much more energy than simply cooling air, so that systems creating condensation may use a large amount of electricity or other energy. Improvements in operations of the system may occur because, if pipes in the system carry water that is below the saturation temperature of the air around the pipes, condensation might form on the pipes. That condensation can damage the pipes or equipment in the conditioned space, cause mold, and cause water to pool on the floor, and can require the installation of insulation on the pipes (to stop the condensation).

In the system shown in FIG. 1, use of elevated temperatures may substantially reduce, or almost entirely eliminate, the need for energy-intensive cooling components such as chillers and the like, even where the heat load in the data center **101** is very high. As a result, system **100** may be operated at a lower operating cost than could otherwise be achieved. In addition, lower capital costs may be required,

6

because fans, coils, heat exchangers, and cooling towers are relatively basic and inexpensive components. In addition, by operating with a higher temperature difference between cooled air and cooling water, less volume of cooling water is needed, thus reducing the size and cost of piping, and the cost to operate pumps and other such components.

In addition, those components are often very standardized, so that their acquisition costs are lower, and they are more easily located, particularly in developing countries and remote areas where it may be beneficial to place a data center **101**. Use of system **100** in remote areas and other areas with limited access to electrical power is also helped by the fact that system **100** may be operated using less electrical power. As a result, such a system can be located near lower-power electrical sub-stations and the like. As discussed more completely below, lower-powered systems may also be amenable to being implemented as self-powered systems using energy sources such as solar, wind, natural-gas powered turbines, fuel cells, and the like.

A data center **101** in sectional view is a building that houses a large number of computers or similar heat-generating electronic components. A workspace **106** is defined around the computers, which are arranged in a number of parallel rows and mounted in vertical racks, such as racks **102a**, **102b**. The racks may include pairs of vertical rails to which are attached paired mounting brackets (not shown). Trays containing computers, such as standard circuit boards in the form of motherboards, may be placed on the mounting brackets.

In one example, the mounting brackets may be angled rails welded or otherwise adhered to vertical rails in the frame of a rack, and trays may include motherboards that are slid into place on top of the brackets, similar to the manner in which food trays are slid onto storage racks in a cafeteria, or bread trays are slid into bread racks. The trays may be spaced closely together to maximize the number of trays in a data center, but sufficiently far apart to contain all the components on the trays and to permit air circulation between the trays.

Other arrangements may also be used. For example, trays may be mounted vertically in groups, such as in the form of computer blades. The trays may simply rest in a rack and be electrically connected after they are slid into place, or they may be provided with mechanisms, such as electrical traces along one edge, that create electrical and data connections when they are slid into place.

Air may circulate from workspace **106** across the trays and into warm-air plenums **104a**, **104b** behind the trays. The air may be drawn into the trays by fans mounted at the back of the trays (not shown). The fans may be programmed or otherwise configured to maintain a set exhaust temperature for the air into the warm air plenum, and may also be programmed or otherwise configured to maintain a particular temperature rise across the trays. Where the temperature of the air in the work space **106** is known, controlling the exhaust temperature also indirectly controls the temperature rise. The work space **106** may, in certain circumstances, be referenced as a “cold aisle,” and the plenums **104a**, **104b** as “warm aisles.”

The temperature rise can be large. For example, the work space **106** temperature may be about 77° F. (25° C.) and the exhaust temperature into the warm-air plenums **104a**, **104b** may be set to 113° F. (45° C.), for a 36° F. (20° C.) rise in temperature. The exhaust temperature may also be as much as 212° F. (100° C.) where the heat generating equipment can operate at such elevated temperature. For example, the temperature of the air exiting the equipment and entering the

warm-air plenum may be 118.4, 122, 129.2, 136.4, 143.6, 150.8, 158, 165, 172.4, 179.6, 186.8, 194, 201, or 208.4° F. (48, 50, 54, 58, 62, 66, 70, 74, 78, 82, 86, 90, 94, or 98° C.). Such a high exhaust temperature generally runs contrary to teachings that cooling of heat-generating electronic equipment is best conducted by washing the equipment with large amounts of fast-moving, cool air. Such a cool-air approach does cool the equipment, but it also uses lots of energy.

Cooling of particular electronic equipment, such as microprocessors, may be improved even where the flow of air across the trays is slow, by attaching impingement fans to the tops of the microprocessors or other particularly warm components, or by providing heat pipes and related heat exchangers for such components.

The heated air may be routed upward into a ceiling area, or attic 105, or into a raised floor or basement, or other appropriate space, and may be gathered there by air handling units that include, for example, fan 110, which may include, for example, one or more centrifugal fans appropriately sized for the task. The fan 110 may then deliver the air back into a plenum 108 located adjacent to the workspace 106. The plenum 108 may be simply a bay-sized area in the middle of a row of racks, that has been left empty of racks, and that has been isolated from any warm-air plenums on either side of it, and from cold-air work space 106 on its other sides. Alternatively, air may be cooled by coils defining a border of warm-air plenums 104a, 104b and expelled directly into workspace 106, such as at the tops of warm-air plenums 104a, 104b.

Cooling coils 112a, 112b may be located on opposed sides of the plenum approximately flush with the fronts of the racks. (The racks in the same row as the plenum 108, coming in and out of the page in the figure, are not shown.) The coils may have a large surface area and be very thin so as to present a low pressure drop to the system 100. In this way, slower, smaller, and quieter fans may be used to drive air through the system. Protective structures such as louvers or wire mesh (not shown) may be placed in front of the coils 112a, 112b to prevent them from being damaged.

In operation, fan 110 pushes air down into plenum 108, causing increased pressure in plenum 108 to push air out through cooling coils 112a, 112b. As the air passes through the coils 112a, 112b, its heat is transferred into the water in the coils 112a, 112b, and the air is cooled.

The speed of the fan 110 and/or the flow rate or temperature of cooling water flowing in the cooling coils 112a, 112b may be controlled in response to measured values. For example, the pumps driving the cooling liquid may be variable speed pumps that are controlled to maintain a particular temperature in work space 106. Such control mechanisms may be used to maintain a constant temperature in workspace 106 or plenums 104a, 104b and attic 105.

The workspace 106 air may then be drawn into racks 102a, 102b such as by fans mounted on the many trays that are mounted in racks 102a, 102b. This air may be heated as it passes over the trays and through power supplies running the computers on the trays, and may then enter the warm-air plenums 104a, 104b. Each tray may have its own power supply and fan, with the power supply at the back edge of the tray, and the fan attached to the back of the power supply. All of the fans may be configured or programmed to deliver air at a single common temperature, such as at a set 113° F. (45° C.). The process may then be continuously readjusted as fan 110 captures and circulates the warm air.

Additional items may also be cooled using system 100. For example, room 116 is provided with a self-contained fan-coil unit 114 which contains a fan and a cooling coil. The

unit 114 may operate, for example, in response to a thermostat provided in room 116. Room 116 may be, for example, an office or other workspace ancillary to the main portions of the data center 101.

In addition, supplemental cooling may also be provided to room 116 if necessary. For example, a standard roof-top or similar air-conditioning unit (not shown) may be installed to provide particular cooling needs on a spot basis. As one example, system 100 may be designed to deliver 78° F. (25.56° C.) supply air to work space 106, and workers may prefer to have an office in room 116 that is cooler. Thus, a dedicated air-conditioning unit may be provided for the office. This unit may be operated relatively efficiently, however, where its coverage is limited to a relatively small area of a building or a relatively small part of the heat load from a building. Also, cooling units, such as chillers, may provide for supplemental cooling, though their size may be reduced substantially compared to if they were used to provide substantial cooling for the system 100.

Fresh air may be provided to the workspace 106 by various mechanisms. For example, a supplemental air-conditioning unit (not shown), such as a standard roof-top unit may be provided to supply necessary exchanges of outside air. Also, such a unit may serve to dehumidify the workspace 106 for the limited latent loads in the system 100, such as human perspiration. Alternatively, louvers (not shown) may be provided from the outside environment to the system 100, such as powered louvers to connect to the warm air plenum 104b. System 100 may be controlled to draw air through the plenums when environmental (outside) ambient humidity and temperature are sufficiently low to permit cooling with outside air. Such louvers may also be ducted to fan 110, and warm air in plenums 104a, 104b may simply be exhausted to atmosphere, so that the outside air does not mix with, and get diluted by, the warm air from the computers. Appropriate filtration may also be provided in the system, particularly where outside air is used.

Also, the workspace 106 may include heat loads other than the trays, such as from people in the space and lighting. Where the volume of air passing through the various racks is very high and picks up a very large thermal load from multiple computers, the small additional load from other sources may be negligible, apart from perhaps a small latent heat load caused by workers, which may be removed by a smaller auxiliary air conditioning unit as described above.

Cooling water may be provided from a cooling water circuit powered by pump 124. The cooling water circuit may be formed as a direct-return, or indirect-return, circuit, and may generally be a closed-loop system. Pump 124 may take any appropriate form, such as a standard centrifugal pump. Heat exchanger 122 may remove heat from the cooling water in the circuit. Heat exchanger 122 may take any appropriate form, such as a plate-and-frame heat exchanger or a shell-and-tube heat exchanger.

Heat may be passed from the cooling water circuit to a condenser water circuit that includes heat exchanger 122, pump 120, and cooling tower 118. Pump 120 may also take any appropriate form, such as a centrifugal pump. Cooling tower 118 may be, for example, one or more forced draft towers or induced draft towers. The cooling tower 118 may be considered a free cooling source, because it requires power only for movement of the water in the system and in some implementations the powering of a fan to cause evaporation; it does not require operation of a compressor in a chiller or similar structure.

The cooling tower 118 may take a variety of forms, including as a hybrid cooling tower. Such a tower may

combine both the evaporative cooling structures of a cooling tower with a water-to-water heat exchanger. As a result, such a tower may be fit in a smaller space and be operated more modularly than a standard cooling tower with separate heat exchanger. An additional advantage may be that hybrid towers may be run dry, as discussed above. In addition, hybrid towers may also better avoid the creation of water plumes that may be viewed negatively by neighbors of a facility.

As shown, the fluid circuits may create an indirect water-side economizer arrangement. This arrangement may be relatively energy efficient, in that the only energy needed to power it is the energy for operating several pumps and fans. In addition, this system may be relatively inexpensive to implement, because pumps, fans, cooling towers, and heat exchangers are relatively technologically simple structures that are widely available in many forms. In addition, because the structures are relatively simple, repairs and maintenance may be less expensive and easier to complete. Such repairs may be possible without the need for technicians with highly specialized knowledge.

Alternatively, direct free cooling may be employed, such as by eliminating heat exchanger 122, and routing cooling tower water (condenser water) directly to cooling coils 112a, 112b, as shown in FIG. 2. Such an implementation may be more efficient, as it removes one heat exchanging step. However, such an implementation also causes water from the cooling tower 118 to be introduced into what would otherwise be a closed system. As a result, the system in such an implementation may be filled with water that may contain bacteria, algae, and atmospheric contaminants, such as leaves, dirt, bird feathers or excrement, and insects or other animals that may enter the tank or mold and may also be filled with other contaminants in the water. A hybrid tower, as discussed above, may provide similar benefits without the same detriments.

In order to keep the water in the cooling tower 118 clean, a water treatment facility or system 205 can treat water that flows into and out of the cooling tower 118. Because the cooling tower 118 is open to the environment, the tower is prone to collecting any debris that is in the environment or other items that may accumulate and potentially clog the water transfer system, i.e., the pipes. The water treatment system 205 can reduce or eliminate these contaminants.

The complexity of the water treatment system 205 can vary from a very simple filtration system to a full-scale water treatment system on par with municipal water purification or sewage treatment facilities. The system can remove physical, chemical or biological contaminants or any combination thereof. A very simple system may include a filter-based system, which filters out contaminants over a particular size threshold, such as large contaminants, e.g., anything over 1 inch in size, or contaminants with a much smaller size, e.g., anything over 25 microns in size. Another simple system can include settling tanks, which allow contaminants to settle out of the water. A flocculation tank can be used, which aggregates smaller particles into larger particles, which then can be more easily removed from the water. Septic tanks, sedimentation tanks, fluidized bed reactors, fixed film system or suspended growth systems, activated sludge systems, aeration tanks, biofilters or aerobic treatment systems can also be included in the waste treatment system. Cooling water can be treated by ionization, UV treatment, reverse osmosis, or chemicals, such as biocides or ozone, to form highly clean water which will then reduce the need for maintenance on the system because growth of organisms can be prevented in the super cleaned water. The water treatment

apparatus can include any combination of these components or methods of treatment, in addition to a single method by itself.

Determining how to configure the waste treatment system can include determining the desired purity level of the treated water, determining the acceptable contaminant size of the treated water, as well as the likely contaminant size of the incoming water and an amount of space, e.g., square footage or acreage, available or required to house the water treatment system. Additionally, the cost of treatment can be a factor. The cost of purifying the water to potable quality can be sufficiently high that it may not be economically feasible (or may be more expensive than the cost of repairing or replacing piping in the facility caused by contaminants), particularly because it will become dirty from exposure to the atmosphere once again. On the other hand, the cost of maintenance can be reduced if the water is sufficiently cleaned to prevent the system from becoming damaged by contaminants. A suitable type of water treatment apparatus can be determined based on the desire to expend resources on maintaining the system versus the cost of treating the water using the selected apparatus. In some embodiments, the water is only partially treated and is pumped back into the system along with a source of cleaner water. This water is referred to herein as brown water, not because it necessarily has been used in household applications, but because it has already been used in the cooling system. Inputting brown water, after only partial treatment, can be used to reduce the need for clean water from a municipal source.

In addition or as an alternative to adding a water treatment system 205, a filter may be placed in line (not shown) between the cooling tower 118 and the data center 101, such as on an outlet of the cooling tower 118. This filter can prevent large contaminants from entering the pipes that direct water from the cooling tower 118 to the data center 101. The filter can be a single filter or a series of filters for capturing contaminants and particulates.

The water in the cooling tower 118 can be flowed to the treatment system 205 continuously or at regular intervals, so that the cooling tower 118 remains at least partially full at any time during operation of the data center 101. Because the cooling tower 118 can accumulate sediment or sludge on the interior sidewalls, the entire tower 118 may be periodically emptied to clean the walls, such as by scrubbing out the interior of the tower. The location of the cooling tower and seasons may dictate how frequently the tower requires cleaning. For example, bacteria or algae may be more active in warmer months than in cooler months and therefore more frequent cleaning or treatment may be required during the summer months.

In some embodiments, a portion of the water and contaminants are removed from the waste treatment system 205 and instead of being returned to the cooling tower 118, are released as waste, such as to a county or municipal sewer system. The percentage that is released as waste can depend on the acceptable limits set by the county or municipality regarding how concentrated an effluent waste stream from a facility can be. For example, the waste that is released from the waste treatment system 205 may be a combination of sludge and liquid. Release of some percentage of waste from the water system allows for flushing minerals out of the system that might otherwise build up over time and potentially damage the system.

The waste treatment system 205 is not a part of a governmental water pollution control facility. Specifically, the waste treatment system 205 is a part of the system 100, rather than a separate facility. The waste treatment system

11

205 may be physically proximate to the data center 101, such as on a common piece of property, that is, property that is all part of a single tract, is commonly owned, is within a defined space, such as within a tract of land suitably sized for an industrial use, such as 50 acres, or is commonly fenced in. Because the waste treatment system 205 is near the data center 101 and the cooling system (e.g., on the same site or an adjacent site), water can be returned from the waste treatment system 205 directly to the cooling system.

In addition to returning water to the cooling system, such as the cooling tower 118, from the waste treatment system 205, a source of fresh water can also be introduced into the cooling tower. Control valve 126 is provided in the condenser water circuit to supply make-up water to the circuit. Make-up water may generally be needed because cooling tower 118 operates by evaporating large amounts of water from the circuit. The control valve 126 may be tied to a water level sensor in cooling tower 118, or to a basin shared by multiple cooling towers. When the water falls below a predetermined level, control valve 126 may be caused to open and supply additional makeup water to the circuit. A back-flow preventer (BFP) may also be provided in the make-up water line to prevent flow of water back from cooling tower 118 to a main water system, which may cause contamination of such a water system.

In some embodiments, a 30-40 MW facility has a peak flow of about 10,000 gallons/minute and uses up to about 15 million gallons per day, e.g., 10 million gallons, for cooling. Approximately 5-10% of the water flowing through the system can be lost to evaporation, depending on the dimensions of the cooling tower and the ambient conditions. For example, a narrow and deep cooling tower can reduce loss of water in the cooling tower, such as to entrainment in the air stream over the cooling tower.

Approximately 1% of the flow through the data center can be directed to the waste treatment system 205. Some amount of the flow into the waste treatment system will be released off-site, such as to the sewer, while the remainder is returned to the cooling tower. The amount of water that is treated from each cycle of water through the cooling system can vary depending on the efficiency of the water treatment system, the tolerance for contaminants in the cooling system and acceptable cost of treatment. Achieving a higher cleanliness of treated water comes at higher capital cost, but reduces water consumption by reducing the amount of make up water needed by the cooling system. In some embodiments, a larger percentage of water is removed from the system for on-site treatment, such as 10%, each day. A percentage of the treated water, the waste percentage, can then be sent off-site for disposal or further treatment. The more water that is removed from the system for cleaning each day, the lower the level of water cleanliness coming out of the waste treatment system back to the cooling tower needs to be to achieve sufficiently clean water that is acceptable for the desired maintenance schedule.

A portion of the treated water is returned to the cooling tower for reuse. Because the water is used for cooling and not for human consumption, the quality of the treated water may not be that of water sourced from a potable drinking source, but can be sufficiently clean to prevent damage to the cooling system. Because there will be some loss of water either to evaporation or as waste that is sent off-site, the make-up water is used to ensure that there is sufficient water in the system to keep the data center sufficiently cooled or that there is sufficient water to keep the total amount of water in the system constant. The make-up water can be provided by a main water system, e.g., a municipal water system, or

12

from another source, such as captured precipitation, surface water or ground water. In some systems, the source of make-up water is effluent from an industrial or agricultural plant that either is sufficiently clean for delivering to the system or which can be cleaned by the waste treatment system for delivering to the system. Because many plants are required to have their own treatment centers on-site, the water coming out of these plants may be clean enough to use with the cooling system. Further, diverting this effluent to the data center cooling system can allow the plant to save on sewer costs.

Referring back to FIG. 1, optionally, a separate chiller circuit may be provided. Operation of system 100 may switch partially or entirely to this circuit during times of extreme atmospheric ambient (i.e., hot and humid) conditions or times of high heat load in the data center 101. Controlled mixing valves 134 are provided for electronically switching to the chiller circuit, or for blending cooling from the chiller circuit with cooling from the condenser circuit. Pump 128 may supply tower water to chiller 130, and pump 132 may supply chilled water, or cooling water, from chiller 130 to the remainder of system 100. Chiller 130 may take any appropriate form, such as a centrifugal, reciprocating, or screw chiller, or an absorption chiller.

The chiller circuit may be controlled to provide various appropriate temperatures for cooling water. In some implementations, the chilled water may be supplied exclusively to a cooling coil, while in others, the chilled water may be mixed, or blended, with water from heat exchanger 122, with common return water from a cooling coil to both structures. The chilled water may be supplied from chiller 130 at temperatures elevated from typical chilled water temperatures. For example, the chilled water may be supplied at temperatures of 55° F. (13° C.) to 65 to 70° F. (18 to 21° C.) or higher. The water may then be returned at temperatures like those discussed below, such as 59 to 176° F. (15 to 80° C.). In this approach that uses sources in addition to, or as an alternative to, free cooling, increases in the supply temperature of the chilled water can also result in substantial efficiency improvements for the system 100.

Pumps 120, 124, 128, 132, may be provided with variable speed drives. Such drives may be electronically controlled by a central control system to change the amount of water pumped by each pump in response to changing set points or changing conditions in the system 100. For example, pump 124 may be controlled to maintain a particular temperature in workspace 106, such as in response to signals from a thermostat or other sensor in workspace 106.

In operation, system 100 may respond to signals from various sensors placed in the system 100. The sensors may include, for example, thermostats, humidistats, flowmeters, and other similar sensors. In one implementation, one or more thermostats may be provided in warm air plenums 104a, 104b, and one or more thermostats may be placed in workspace 106. In addition, air pressure sensors may be located in workspace 106, and in warm air plenums 104a, 104b. The thermostats may be used to control the speed of associated pumps, so that if temperature begins to rise, the pumps turn faster to provide additional cooling waters. Thermostats may also be used to control the speed of various items such as fan 110 to maintain a set pressure differential between two spaces, such as attic 105 and workspace 106, and to thereby maintain a consistent airflow rate. Where mechanisms for increasing cooling, such as speeding the operation of pumps, are no longer capable of keeping up with increasing loads, a control system may activate chiller

13

130 and associated pumps 128, 132, and may modulate control valves 134 accordingly to provide additional cooling.

Various values for temperature of the fluids in system 100 may be used in the operation of system 100. In one exemplary implementation, the temperature setpoint in warm air plenums 104a, 104b may be selected to be at or near a maximum exit temperature for trays in racks 102a, 102b. This maximum temperature may be selected, for example, to be a known failure temperature or a maximum specified operating temperature for components in the trays, or may be a specified amount below such a known failure or specified operating temperature. In certain implementations, a temperature of 45° C. may be selected. In other implementations, temperatures of 25° C. to 125° C. may be selected. Higher temperatures may be particularly appropriate where alternative materials are used in the components of the computers in the data center, such as high temperature gate oxides and the like.

In one implementation, supply temperatures for cooling water may be 68° F. (20° C.), while return temperatures may be 104° F. (40° C.). In other implementations, temperatures of 50° F. to 84.20° F. or 104° F. (10° C. to 29° C. or 40° C.) may be selected for supply water, and 59° F. to 176° F. (15° C. to 80° C.) for return water. Chilled water temperatures may be produced at much lower levels according to the specifications for the particular selected chiller. Cooling tower water supply temperatures may be generally slightly above the wet bulb temperature under ambient atmospheric conditions, while cooling tower return water temperatures will depend on the operation of the system 100.

Using these parameters and the parameters discussed above for entering and exiting air, relatively narrow approach temperatures may be achieved with the system 100. The approach temperature, in this example, is the difference in temperature between the air leaving a coil and the water entering a coil. The approach temperature will always be positive because the water entering the coil is the coldest water, and will start warming up as it travels through the coil. As a result, the water may be appreciably warmer by the time it exits the coil, and as a result, air passing through the coil near the water's exit point will be warmer than air passing through the coil at the water's entrance point. Because even the most-cooled exiting air, at the cooling water's entrance point, will be warmer than the entering water, the overall exiting air temperature will need to be at least somewhat warmer than the entering cooling water temperature.

Keeping the approach temperature small permits a system to be run on free, or evaporative, cooling for a larger portion of the year and reduces the size of a needed chiller, if any is needed at all. To lower the approach temperature, the cooling coils may be designed for counterflow rather than for self-draining. In counter-flow, the warmest air flows near the warmest water and the coolest air exits near where the coolest water enters.

In certain implementations, the entering water temperature may be 64° F. (18° C.) and the exiting air temperature 76.6° F. (24.8° C.), as noted above, for an approach temperature of 12.6° F. (10.8° C.). In other implementations, wider or narrower approach temperature may be selected based on economic considerations for an overall facility.

With a close approach temperature, the temperature of the cooled air exiting the coil will closely track the temperature of the cooling water entering the coil. As a result, the air temperature can be maintained, generally regardless of load, by maintaining a constant water temperature. In an evapo-

14

orative cooling mode, a constant water temperature may be maintained as the wet bulb temperature stays constant (or changes very slowly), and by blending warmer return water with supply water as the wet bulb temperature falls. As such, active control of the cooling air temperature can be avoided in certain situations, and control may occur simply on the cooling water return and supply temperatures. The air temperature may also be used as a check on the water temperature, where the water temperature is the relevant control parameter.

FIG. 3A is a psychrometric chart showing a heating and cooling cycle for air in a data center. A psychrometric chart graphically represents the thermodynamic properties of moist air (which is air containing any appreciable moisture, and not merely air that would feel moist to a person). The chart is from ASHRAE Psychrometric Chart No. 1, which defines properties of air for sea level applications. See 1997 ASHRAE Handbook—Fundamentals, at page 6.15. Other charts may also be used, and the chart shown here is merely used to exemplify certain aspects of the concepts discussed in the present document.

The psychrometric chart is criss-crossed with a number of lines that represent various properties of air. Cooling and heating processes on air can be analyzed by identifying a point on the chart that represents air at a particular condition (e.g., temperature and humidity), and then locating another point that represents the air at another condition. A line between those points, generally drawn as a straight line, may be fairly assumed to represent the conditions of the air as it moves from the first condition to the second, such as by a cooling process.

Several properties will be discussed here. First, saturation temperature is an arc along the left of the graph and represents the temperature at which air becomes saturated and moisture begins coming out of the air as a liquid—also known popularly as the “dew point.” When the temperature of air is taken below the dew point, more and more water comes out of the air because the cooler air is capable of holding less water.

The dry bulb temperature of air is listed along the bottom of the graph and represents what is popularly viewed as temperature, i.e., the temperature returned by a typical mercury thermometer.

The chart shows two numbers relating to the humidity of air. The first is the humidity ratio, listed along the right edge of the graph, and is simply the weight of moisture per each unit weight of dry air. Thus, the humidity ratio will stay constant at various temperatures of air, until moisture is removed from the air, such as by pushing the air temperature down to its dew point (e.g., the moisture comes out of the air and ends up on the grass in the morning) or by putting moisture into the air (e.g., by atomizing water into such as fine mist in a humidifier that the mist can be supported by the natural motion of the air molecules). Thus, when graphing processes that involve simple changes in air temperature, the point that represents the state of the air will move straight left and right along the graph at a constant humidity ratio. That is because the dry-bulb temperature will go up and down, but the humidity ratio will stay constant.

The second humidity parameter is the so-called relative humidity. Unlike the humidity ratio, which measures the absolute amount of moisture in the air, the relative humidity measures the amount of moisture in the air as a percent of the total moisture the air could possibly hold at its current temperature. Warmer air can hold more moisture than can colder air, because the molecules in warmer air are moving more rapidly. Thus, for an equal amount of moisture in the

15

air (i.e., an equal humidity ratio), the relative humidity will be lower at a high temperature than at a low temperature.

As one example, on a summer day when the overnight low was 55 degrees Fahrenheit and there is dew on the ground, and the daytime high is around 75 degrees Fahrenheit, the relative humidity in the early morning will be about 100% (the dew point), but the relative humidity in the afternoon will be a very comfortable 50%, even if one assumes that the identical amount of water is in the air at both times. This exemplary process is shown in FIG. 3 by the points marked C and D, with point C showing saturated air at 55 degrees Fahrenheit (the overnight low), and point D representing that same air warmed to 75 degrees Fahrenheit (the daytime temperature).

Commercial air handling systems take advantage of this same process in providing conditioned air in a building. Specifically, systems may gather air from an office space at 75 degrees Fahrenheit and a relative humidity of 60 percent. The systems pass the air through a cooling coil that looks like an automobile radiator to cool the air to 55 degrees Fahrenheit, which will typically push the air down to its dew point. This will make moisture pour out of the air as it passes through the cooling coil. The moisture can be captured in drains below the cooling coil and then be removed from the building. The air can then be returned to the work space, and when it warms back to 75 degrees, it will be a very comfortable 50 percent relative humidity.

This common cooling process is shown by points A, B, C, and D on the chart of FIG. 3A. Point A shows the 75 degree air at 70 percent relative humidity. Point B shows the air cooled to its dew point, which it hits at a temperature (dry bulb) of about 65 degree Fahrenheit. Further cooling of the air to 55 degrees Fahrenheit (to point C) rides along the saturation curve, and water will come out of the air during that portion of the cooling. Finally, the state of the air moves to point D as the air is warmed and reaches 75 degrees Fahrenheit again. At this point, the relative humidity will be 50 percent (assuming it does not pick up additional moisture from the room or the existing room air) rather than the original 60 percent because the cooling process has dehumidified the air by pulling moisture out of the air in the cooling coil. If the room air contains more moisture than does the cooled air, point D will be slightly above its position shown in FIG. 2, but still below point A.

Such a common process brings with it a number of challenges. First, to cool the air to 55 degrees Fahrenheit, the system must provide cooling water in the cooling coil that can absorb all the heat. Such water would need to be at least cooler than 55 degrees Fahrenheit. It may be expensive to create such cool water—requiring systems such as chillers and other energy-intensive systems. In addition, the area immediately around the pipes that supply the cooling water will be cooler than 55 degrees Fahrenheit, i.e., cooler than the dew point of the air if the pipes run through the workspace or through air having the same state as the air in the workspace. As a result, moisture from the air may condense on the pipe because the temperature of the surrounding air has fallen to its dew point. Thus, insulation may be required around the cool pipes to prevent such condensation, and condensation might occur in any event, and cause rusting, mold, water pooling, or other problems. Finally, as any high school physics student likely knows, it takes a lot of energy to dehumidify, i.e., to change the water from one state to another.

The warm air cooling features discussed with respect to FIG. 1 above may, in certain implementations, avoid one or more of these challenges. An exemplary warm air cooling

16

process is shown on the graph of FIG. 2 by points E and F. Point E shows a room-air condition in a workspace that is near the top of, but within, common guidelines for comfort levels for people dressed in summer clothing. See 1997 ASHRAE Handbook—Fundamentals, at page 8.12. That condition is 75° F. (24° C.), and a relative humidity of about 70 percent (the same as Point A in the prior example). Point F shows heating of that air without the addition of moisture, such as by passing the air over heat-generating computer components in a rack-mounted server system. The temperature rise is 36° F. (2° C.) to bring the air to a state of 111° F. (44° C.) at about 23 percent relative humidity. The air may then be cooled to its original temperature (point E) of 75 degrees Fahrenheit in a cooling coil before it is re-introduced to the work space, without adding water to or removing water from the air.

Points G and H on the graph represent the condition of air in the space immediately surrounding cooling pipes. It is assumed for this example that the cooling supply water is 68 degrees Fahrenheit (20 degrees Celsius) and the return temperature is 104 degrees Fahrenheit (40 degrees Celsius). It is also assumed that the air near the pipes will contain the same moisture level as the rest of the air in the space, and that the air immediately surrounding the pipe takes on the same temperature as the water inside the pipe. As can be seen, this air associate with the cooling water also stays above the saturation point, so that there should be no condensation on the cooling water pipes, and thus no need for insulation to prevent condensation on the pipes.

It can be seen by this process that the air never becomes saturated. As a result, the system need not provide energy to create a phase change in the air. In addition, the system need not provide liquid recovery structures at the cooling coil, or pipe insulation for anywhere. Other similar temperatures, and in many implementations warmer temperatures, may be used. The particular temperatures discussed here are meant to be exemplary only.

FIG. 3B is a graph of setpoint temperature for a computing facility over a one year time period. The setpoint temperature may be a temperature in a work space such as work space 106 in FIG. 1. An example outdoor wetbulb temperature through the year is also shown for comparison. As shown in the stepped graph, the setpoint temperature (i.e., a targeted temperature) is adjusted infrequently, such as with the seasons or monthly, so as to more closely track expected outdoor temperatures. The setpoint is increased in the summer because the lower winter setpoint may be effectively unattainable in warm summer weather by using only evaporative cooling. Thus, while a “best,” lower, setpoint temperature applies in the winter, that same setpoint is not realistic, in the example, in summer months.

The setpoint is adjusted infrequently so as to better approximate a setpoint that is attainable using evaporative cooling techniques with little or no assistance from chillers or other similar components that require relatively high levels of energy to operate. Although increasing the setpoint during warmer times of the year may increase the typical operating temperature, it also decreases the amount of thermal cycling that may occur in a facility, and thus lengthen the life of electronic components in the facility. In contrast, if the setpoint is kept as low as possible, the conditioned space would be relative cool on days having a low wet bulb temperature and relatively warm on days having a high wet bulb temperature. Thus, keeping a constant setpoint throughout the year may actually increase

thermal cycling, particularly in warmer months—as the system is able to maintain the setpoint on some days but not on other days.

The setpoint may also be adjusted substantially continuously, such as by varying the setpoint temperature in an annual sinusoidal manner that generally follows the expected outdoor wet bulb temperature, as shown by the continuous setpoint line. Studies have indicated that human discomfort is minimized by providing many minute changes, or continuous changes, to temperature, as opposed to large step variations in temperature. In both examples the setpoint may be maintained, in certain implementations, even if a lower temperature may readily be achieved (e.g., because the outdoor wet bulb temperature is lower than expected) so as to minimize thermal cycling in a facility being cooled.

The particular setpoint temperatures may be selected based on the capabilities of the components in a facility and on prevailing local weather conditions. For example, cool weather setpoints may be in the range of 59-77° F. (15-25° C.), with particular values of 64.4, 68, 71.6, and 75.2° F. (18, 20, 22, and 24° C.). Warm weather setpoints may be in the range of 68-86° F. (20-30° C.), with particular values of 71.6, 75.2, 78.8, and 82.4° F. (22, 24, 26, and 28° C.). In a particular implementation, warm weather air temperatures in a facility may be approximately 80.60-82.40° F. (27-28° C.) and cold weather temperature may be about 71.60° F. (22° C.). The time for resetting the setpoint may also vary, and may be weekly (e.g., using a long range weather forecast to select an achievable setpoint that tracks the predicted wet bulb temperature), weekly, monthly, or quarterly, for example.

When the wet bulb temperature gets too high to achieve the desired set point, the temperature of the cooling water may be allowed to drift upward with the wet bulb temperature, causing the temperature in space 106 to move upward also. Alternatively, cooling such as from chiller 130 may be provided, and the chilled water may be blended with other cooling water so that the setpoint may be maintained.

Cooling with relatively warm water may also provide certain benefits when chillers are used. In particular, when a chiller is allowed to provide a smaller temperature change to a coolant, it may provide cooling for less electrical consumption per ton of cooling than if it were required to impart a greater temperature change to the coolant. By having elevated air temperatures in a cooled space, the supply water temperature may likewise be higher, and the need for a chiller to cool the water may be less.

FIG. 4 is a plan view in schematic form of a hydronic cooling system for a data center 300. The data center 300 includes an existing plant 302, and a future plant 304. As will be seen, the system is highly modular, and can therefore be scaled from a very small system to a very large system by adding additional components and subsystems. Although the existing and future systems are shown as separate systems, extensions of the existing system may also be employed for future expansion. For example, oversized piping may be installed in some implementations and may be tapped for future additions of added components when the data center 300 is expanded. Also, other forms of cooling loops, or additional or alternative cooling loops may be employed.

The cooling system for existing plant 302 is similar to that shown in FIG. 1. Cold air is supplied to a workspace 306 by air handling unit 310. Air handling unit 310 may include a supply fan (not shown) that draws warm air out of a plenum, such as in an above-ceiling attic space, and pushes the air through cooling coils 312, 314. The warm air may be

produced by passing cooled air over computers, such as computer, mounted on trays in a rack system. Air passing over the computers may be collected in warm-air plenums 316, and may be routed to the above-ceiling space or another appropriate space.

In alternative implementations, multiple air handling units that include a fan and a cooling coil may be placed immediately above or below the warm air plenums 316. Such air handling units may take the form of long banks of small centrifugal fans and long heat exchange units mounted to the fans. The fans may pull air or push air through the coils. Pushing air may have the advantage of being quieter, as the coils may block out a certain amount of the fan noise. Also pushing of air may be more efficient. Pulling of air may provide a benefit of allowing a limited number of fans to operate on a much larger bank of coils, as all the pulling fans can be connected to a plenum, and may create a relative vacuum behind the coils to pull air through. In such an arrangement, if one of the fans breaks down, the others can more easily provide support across the entire coil length.

Cooling water may be provided by pumps 326 to cooling coils 312, 314 to cool the air. These pumps 326 may draw cooled water from heat exchangers 324, and drive the water into existing plant 302, through cooling coils 312, 314, and back through heat exchangers 324, where the heat acquired from coils 312, 314, may be removed from the water.

Heat in the cooling water may be removed from heat exchangers 324 by cooled tower water supplied through pump 322 from cooling towers 320. The system is shown for exemplary purposes as having two parallel circuits for cooling water and one circuit for tower water from the cooling towers 320. However other appropriate arrangements for the piping system may also be used. For example, a single heat exchanger may be used, as may a single cooling tower. Other sources of free cooling may also be used, such as ground water or deep lake cooling.

Also, additional space may be provided for extra cooling towers and future heat exchangers, and piping may be sized to accommodate future expansion. For example, cooling towers may be provided in an extendable stacked arrangement, so that a third tower will be piped in parallel with the first two towers 320 when the additional cooling is needed.

Existing plant 302 includes an additional workspace 308. This space may be for example, an office space, control room, or storage room. The space may be designed to be inhabited extensively by people, and may require lower temperatures than does workspace 306. As a result, workspace 308 is provided with air-conditioning unit 330, which may be a standard rooftop air conditioning unit. Although air-conditioning unit 330 may require additional electricity for providing higher degrees of cooling, including latent heat removal, the unit 330 generally serves an area and a thermal load that is much smaller than the entire system. As a result, unit 330 may be much smaller than what would otherwise be required to provide similar low-temperature cooling for the rest of the system.

Although not shown, one or more chillers and chillers circuits may be provided to operate the system, in addition to the tower cooling water circuit, whether together or alternatively. The piping of the chiller circuit may take a form similar to that shown schematically in FIG. 1.

As noted above, future plant 304 may be provided in a variety of ways. In general, a modular approach may be used, so that additional heat exchangers are added in a size commensurate with the additional loads that are expected. Certain components may be piped as entirely separate circuits, while others may be extensions of the original

system. Sharing components between old and new systems may cause construction-based interruptions in the existing system, but may also provide for better utilization of the components in the full system. For example, condenser water can be shared through the system, so that if each half of the system requires cooling from 1.5 cooling towers, only three towers will be operated. If the two systems are separate, each system would have to run two towers, and thereby use additional energy.

FIG. 5 is a plan view in schematic form of a hydronic cooling system for a data center housed in outdoor shipping containers 402 and 404. The system is arranged in a manner similar to the system in FIG. 4. Here, however, cooling coils are mounted along the length of both of shipping containers 402, 404, such as below or above the computer-holding racks in the system. For example, an elevated false floor may be provided down the center of the shipping containers 402, 404, and squirrel cage or other forms of fans may blow or pull warm air through cooling coils, depositing cool air into the space under the grate. The cooled air may then pass up through the floor grate into cooled-air isles 400, 403, and then drawn by fans associated with each tray in the racks, into warm-air plenums 401, 405 along the sides of the shipping containers, and may be routed back through the fans. Using exterior space along the outer walls for the warm-air plenums 401, 405 may have the advantage of allowing direct rejection of heat on cold days and of minimizing the amount of atmospheric heat that passes into each container on hot days.

In this implementation, the various benefits of using little energy, of modularity, and of standardization on low complexity components may be appreciated. Specifically, shipping containers 402, 404 may be built and outfitted in a central location using standardized techniques and specialized, well-trained labor—a type of prefab construction approach. The shipping containers 402, 404 may then be delivered to a site, such as a site having adequate electrical, water, and data services, and may be easily connected to such services, and then begin their operation.

The other components of the system may be provided separately to the site. For example heat exchangers and pumps may be provided on preassembled skids, dropped into the place, and connected to other components. Cooling towers may also be delivered to a site and connected relatively quickly. Also, because the system may operate without the need for components such as chillers that demand high power levels, the system may be installed in more areas, such as remote areas having lower levels of electrical service.

As shown, the system is also highly modular, so that additional resources may be added relatively quickly and inexpensively as they are needed. Containers 402 are shown as part of an existing system, and containers 404 (with associated equipment) are shown as part of a future addition to the system. Other modules may also be added as the need arises, until the level of available utilities is tapped out. Each container may contain hundreds or thousands of trays, so that a complete installation could include hundreds of thousands of computers.

Systems like those shown in this document may be deployed in a number of different and flexible manners. For example, locations may be selected according to the availability of resources. For example, cooling water may be obtained from ground water or surface water for free cooling processes. In such an environment, cooling towers may be eliminated from the system.

In addition, systems may be deployed far from suitable construction labor (and may use labor having a lower skill level) if they are assembled centrally and shipped to a site.

Such systems may also be self-contained. For example, fuel cells or gas-powered turbines may be deployed with a system to generate electricity, and may be supplemented by wind or solar power. In addition, generators, such as combustion turbines may be selected for less than the full load of a system, but may be deployed to permit peak-shaving of the power required by a system. Such peak-shaving may permit a system to operate in an area that cannot provide all of the power needed for a system. In addition, peak shaving may allow an operator of the system to negotiate energy rates that are lower than would otherwise be available, because the system would not be drawing high levels off the grid when the electrical provider does not have such extra capacity. Moreover, heat from the turbines may be used to help power equipment such as an ammonia-based absorption chiller. In addition, Peltier-like devices may be used to convert heat into electrical energy. In addition, excess pressure from make-up water applied to the cooling towers may also be used to drive small turbines for limited electrical production.

Such systems may also be co-located with other systems that may have use for the hot water generated through the cooling of electronic components. For example, a data center may be located near a hospital or other high density residence to supply heating water, such as for space heating or water heating. In addition, the hot water may be used in common with various processes, such as fermentation, manure digestion, yogurt culture fabrication, and ethanol production. Moreover, the hot water may be used to heat swimming pools or other buildings, or may be run through ice and snow melt piping, such as in airport runways. Likewise, heat from the system may be used to control temperatures in a co-located fish farm. Other implementations may also be employed to serve particular needs for data services, particularly where distribution of processing power closer to end users is needed.

FIG. 6 is a flowchart 600 showing steps for cooling a data center using elevated temperatures. The method is exemplary only; other steps may be added, steps may be removed, and the steps may be performed in different orders than those shown, as appropriate.

At box 602, the exhaust temperatures of various trays in a computer rack system are maintained at a computed value, such as a maximum value for maintaining the trays in reliable operation. The trays may each be self-throttling, in that a fan and related temperature sensor may be associated with each tray. At box 604, at the same time, other components of the system may maintain a workspace temperature in front of the racks at a particular temperature, such as a maximum allowed temperature. The maximum allowed temperature may be, for example, a maximum temperature at which the workers in the workspace are willing to operate. It may also be a temperature set by law or regulation, such as a maximum 40 degree Celsius limit. Federal OSHA and California OSHA guidelines may provide limitations on such temperatures. The set temperature may also change at appropriate times; for example, the temperature may be higher during times when people are not expected to be in the workspace.

At box 606, the workspace cooling is changed in response to changes in the system. Such changes may occur, for example, if the set temperature for the space changes. The changes may also occur, for example, if the temperature remains constant but the load on the system changes—such

21

as if an additional rack of computers is powered up, or if the computers are loaded more heavily. The cooling may change, for example, by increasing the rate of pumping cooling fluid through cooling coils. Alternatively, additional coils may be brought on-line, or a chiller may be started to lower the temperature of cooling water.

At box 608, the exhaust temperature of air from the computer racks is maintained at the maximum computer value. Specifically, although the workspace temperature has changed, and other parameters in the system may have changed, the exhaust temperature may remain unaffected, as each fan associated with a tray remains set on one exhaust temperature.

FIG. 7A shows an alternative embodiment of a cooling system 705 from the system shown in FIG. 1. In addition to a cooling tower 118, the cooling system 705 includes a primary storage tank 710. Water that has been reduced in temperature in the cooling tower 118 is routed to the primary storage tank 710. In some embodiments, the water from the primary water storage tank 710 flows directly from the primary water storage tank 710 to a heat exchanger 122. However, if the water from the primary water storage tank 710 is not sufficiently cool, cooler water can be mixed with the water from the primary water storage tank 710 to bring the temperature down.

In some embodiments, water from the primary storage tank 710 is first routed to a secondary storage unit. The secondary storage unit includes either a single tank or two tanks. Secondary storage units that include two tanks have a buffering tank 720 and a thermal storage tank 722. The buffering tank 720 and thermal storage tank 722 can isolate water within their tanks or may be connected to one another, such as by a series of baffles that allow some water to flow between the two tanks. Water from the primary storage tank 710 flows into, and in some cases, back out of the secondary storage buffering tank 720 to supply the buffering tank 720 with water.

In order to provide cooler water than is available from the primary storage tank 710 for feeding the heat exchanger 122, water from the secondary storage unit is circulated through the evaporator or cold side of the chiller 130. The chiller lowers the temperature of the water in the secondary storage unit to a temperature that is suitable for mixing with water from the primary storage tank 710 for delivery to the heat exchanger 122. If the secondary storage unit includes two compartments, the water can be directed from the buffering tank 720 to the chiller 130, and the thermal storage tank 722 can retain chilled water from the chiller 130 until needed. One advantage of storing chilled water is in the event that the chiller 130 goes off-line and needs to be restarted, a sufficient supply of water can be stored, which allows the cooling system 705 to continue to operate until the chiller 130 is back in service. In some embodiments, a suitable size of a thermal storage tank 722 for each 2 MW of data center power use is about 4000 gallons. A suitable buffering tank 720 size can be about 11,000 gallons.

The condenser side, or hot side, of the chiller 130 can be fed by water from the primary storage tank 710. The chiller then pumps heat from the water taken from the secondary storage unit into water from the primary storage tank 710. The warmed water is then routed to the cooling tower 118 for cooling.

In some embodiments, the primary storage tank 710 holds about the same quantity of water as a small residential swimming pool, such as about 15,000 gallons, which can be suitable for each 2 MW of data center load. At any size, when the temperature drops at night, the water temperature

22

in the primary storage tank 710 and the cooling tower 118 drops as well. This temperature can be sufficiently low for proper operation of the heat exchanger 122 and components (e.g., the coils) in the data center for removing heat from the electronic devices. Thus, the day/night temperature differential can be taken advantage of to reduce use of the chiller 130. The water chilled at night in the primary storage tank 710 can also be transferred to the secondary storage unit for use later in the day when the water temperatures begin to climb. As temperatures climb outside and the temperature in the water from the cooling tower rises, this water is mixed with the stored cooler water from the primary and/or secondary storage tanks so that the water entering the heat exchanger remains in the proper temperature range. This can result in the temperature of the water entering the heat exchanger varying over a 24-hour cycle, being somewhat cooler at night and warming during the day.

The water temperature can also fall so low as to be too cold to feed directly to the heat exchanger 122. As noted herein, water that is too cold can cause condensation if the temperature of the components carrying the chilled water is below that of the dew point. To avoid condensation and the need for expensive insulation, the water through cooling system 705 is kept above a threshold temperature. Thus, it is not desirable to feed water to the heat exchanger 122 that is below that threshold temperature. When the temperature of the water in the primary storage tank 710 is too low, warmed water from the heat exchanger can be fed directly into the primary storage tank 710 (bypassing the cooling tower) to increase the temperature to within the desirable range. Thus, the primary tank acts as a mixer. Alternatively, heated water from the heat exchanger 122 can be mixed with water coming from the primary storage tank 710 and being fed to the heat exchanger 122 (route not shown).

Transferring chilled water from the primary storage tank 710 to the secondary storage unit can provide a source of chilled water without the use of the chiller. As the water in the primary storage tank 710 warms up, chilled water can be drawn directly from the secondary storage unit to reduce the temperature of water supplied to the heat exchanger 122. Thus, variations in ambient temperature, such as day/night variations or day-to-day variations, can be taken advantage of to reduce chiller user.

In some embodiments, multiple primary storage tanks 710 are used with a single data center. With a sufficient number of primary storage tanks 710, a half day or day-long outage of water can be covered by the water in the primary storage tanks. Greater quantities of water can allow for longer outages of water or for shaving off the effects of peak temperatures experienced in the environment over a time greater than a day. However, such quantities of water can be impractical to keep on-site. Under normal operating conditions, the primary storage tanks 710 can be kept full. If a make-up source of water is not available for some period, the primary storage tanks may be allowed to drain to some fraction of their capacity, such as 20% of capacity, until make-up water is again available.

In order to regulate the flow of water to the heat exchanger, whether chilled water is introduced into the flow path, the operation of the chiller and the general flow of water through the cooling system 705, a controller 730 can be included in the cooling system 705. The controller 730 can control valves 134 along each of the water lines as well as operation of the chiller 130 and heat exchanger 122. Optionally, the controller can also provide a warning for data control center that there is a problem with one of the components in the system. The controller can either include

or be in communication with sensors that indicate the ambient temperature and the optimal operating conditions for the system.

As shown in FIG. 7B, the controller can determine the routing of the water through the cooling system according to method 740. The temperature of the water in the primary storage tank can be measured or obtained by the controller (step 742). If the temperature of the water in the primary storage tank is within a suitable range for routing to the heat exchanger, then the water from the primary storage tank is routed directly to the heat exchanger (step 744). If, however, the water is below a suitable range for delivering to the heat exchanger directly, water coming out of the heat exchanger is first mixed with the water from the primary storage tank (step 746). As noted herein, the mixing can occur in the primary storage tank or along the piping that leads to the heat exchanger. The mixed water is then routed to the heat exchanger (step 744).

On the occasions when the water in the primary tank is above the desired temperature range, the system then determines whether there is chilled water stored in the secondary storage tank for mixing with the water from the primary storage tank (step 752). If the water in the secondary tank is sufficiently cool for mixing with the water from the primary tank or if there is a sufficient quantity of chilled water in the secondary storage tank, then water from both the primary and secondary storage tanks are flowed through the cooling system (step 760). The mixed water from the primary and secondary storage tanks is then routed to the heat exchanger (step 744). If the water in the secondary storage tank is not suitably cold or if there is not enough chilled water to mix with the water from the primary storage tank, water can be flowed through the chiller to reduce the temperature (step 764). The water that is chilled by the chiller is then mixed with water from the primary storage tank (step 768). The mixed water is then routed through the heat exchanger (step 744). These steps are then repeated, such as by continuously monitoring the temperature of water flowing through the system and being stored in each of the storage tanks, and adjusting flow as required to keep the water entering the heat exchanger at the desired temperature.

In some implementations of data centers, water returning from the heat exchanger 122 is routed to a holding tank (not shown) beneath the data center. The holding tank can take the place of one of the storage tanks or can be used in addition to the storage tanks. Storing water beneath the data center can reduce piping needs. Alternatively, water from the cooling tower 118 can be routed to a holding tank beneath the data center. Because the water is cooler after cooling off in the cooling tower, there is less likelihood of condensation forming under the data center than if water from the heat exchanger is routed to the holding tank. The holding tank can provide a space efficient means for storing water, particularly if water outages may be possible, such as at particular times of the year when surface water may be frozen and inaccessible.

As in other systems, a water treatment system 205 can be coupled to one or more of the storage tanks or the cooling tower to treat any water and reduce the need for maintenance on the system. Although not shown, the water treatment system 205 can optionally be fed by a non-potable water source to reduce the system's use of municipal or county water. For example, the system 705 can use water that is not sufficiently clean for human consumption and that is from a local source, and treat the water before introducing it to the system 705. Alternatively, a source of make-up water can be fed directly to the primary storage tank 710.

Referring to FIG. 8, in addition to a water treatment apparatus, a settling pond 850 can be located proximate to the cooling system or the water treatment system 205. The settling pond 850 can retain water for sufficiently long periods that any sediment is permitted to settle out of the water. The settling pond 850 can provide for massive thermal storage. Reduced settlement water from the settling pond can also then be routed back to the cooling system. In some embodiments, the water is first treated at the water treatment system 205 before entering the water cooling system. The settling pond 850 can provide further back up water supply when water is not available from a regularly used source, such as city water.

Either the settling pond 850 or another large tank can provide massive thermal storage for the cooling system. For example, a pond or tank that is sufficiently large to hold 100,000 gallons can provide a day's worth of thermal mass for tempering the peaks and troughs of temperature change for each 2 MW of data center load. For seasonal tempering of temperature peaks and troughs, millions of gallons of water may be required for each 2 MW of data center load. Thus, in some embodiments, the settling pond or large storage tank is sized to average out the temperature change of a season or even a year. That is, massive thermal storage can be provided for the cooling system by providing a tank or pond that holds millions of gallons of water.

Optionally, a reflective cover can be located over the settling pond 850. The reflective cover can keep the water in the settling pond 850 from heating up or from evaporating. In some embodiments, the settling pond 850 has a deep central portion where primary settling occurs and a shallow periphery portion provided for wildlife usage. For example, the shallow periphery can include a reed bed or other similar methodology, which can provide a high degree of aerobic biological improvement to the water. The water from the deeper central portion can be recycled to the cooling system, while the water in the shallow periphery is left undisturbed.

Because the data center can be scaled to any size for which sufficient power to run the center is available, the cooling systems can be modularized. In some embodiments, cooling systems are built as modules with a housing that is small enough to be transported over highways, such as, modules having a width less than 13.5 feet, such as about 12 feet, and a length of less than 97 feet, such as less than 30 feet or less than 10 feet. Each module can include one or more of a heat exchanger, a chiller or a controller. Optionally, a module can also include a primary and second water storage tank and a cooling tower, such as on top of the tanks. Alternatively, as shown in FIG. 9, a module with the heat exchanger and controls are connected to a separate system housing the water tanks and the cooling tower.

As shown in FIG. 9, a data center can be constructed with cooling modules 910. In some embodiments, a cooling module 910 includes a heat exchanger, a chiller, a controller, pumps, wiring, valves and piping. The cooling module 910 contains the components in a housing that is up to 13.5 feet wide and 97 feet long, but can be smaller, such as less than 40 feet, less than 30 feet or less than 20 feet long. The cooling module 910 can also be narrower, such as less than about 12 feet, less than 10 feet or less than 8 feet wide. Specifically, the cooling module 910 can be sized to be transportable over roads. The cooling module 910 therefore can be assembled off-site, such as in a factory or other controlled environment, and transported to the data center location. The cooling module 910 is then connected to the data center once in place. Because the cooling module 910

25

is assembled off-site, potentially complex wiring and plumbing jobs can be eliminated from being performed on-site.

The cooling module **910** can be connected to a tank module and cooling tower assembly **920**. The tank module can house the primary water storage tank and the second water storage tank. A cooling tower is then located proximate to the tank module. In some embodiments, a platform is constructed over the tank module and the cooling tower is located on the platform. The primary storage tank and secondary storage tank can each have a capacity of about 15,000 gallons and can have dimensions of about 8 feet in diameter and 40 feet long, that is, in the shape of a long cylindrical tube. Because the tank module and cooling tower assembly **920** is mechanically simple, it can be constructed on-site and connected to the cooling module **910**, such as by piping connected to the cooling module **910**. As the data center grows in size and capacity, more cooling modules **910** and tank module and cooling tower assemblies **920** can be added. This modular system allows for more modules to be added as required.

A number of embodiments have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of what is described. For example, the steps of the exemplary flow chart on FIG. **6** may be performed in other orders, some steps may be removed, and other steps may be added. Accordingly, other embodiments are within the scope of the following claims.

What is claimed is:

1. A system for providing cooled air to electronic devices, comprising:
 - a water storage pond;
 - a reflective cover on the storage pond;
 - a water treatment system fluidly coupled to the storage pond; and
 - an air-to-water heat exchanger connected to the water treatment system and positioned to receive heated air from a group of electronic devices.

26

2. The system of claim **1**, wherein the water treatment system and the air-to-water heat exchanger are on a same site or an adjacent site to one another.

3. The system of claim **2**, further comprising a control system communicably coupled to one or more valves fluidly coupled with the water storage pond, the water treatment system, the air-to-water heat exchanger, and a utility water source.

4. The system of claim **3**, wherein the control system is operable to:

determine that the utility water source is unavailable; and based on the determination, control the one or more valves to route water from one or both of the water storage pond or the water treatment system to the air-to-water heat exchanger.

5. The system of claim **3**, wherein the control system is operable to:

determine that a portion of the water from the water storage pond comprises one or more contaminants; based on the determination, control the one or more valves to route the portion of water from the water storage pond to the water treatment system for treatment; and

control the one or more valves to route the treated water from the water treatment system to the air-to-water heat exchangers.

6. The system of claim **3**, wherein the control system is operable to:

determine a peak cooling power requirement to cool the group of electronic devices based on an ambient environment, the peak cooling power requirement greater than an average cooling power requirement; and based on the determination, control the one or more valves to route water from the water storage pond to meet the peak cooling power requirement.

7. The system of claim **3**, wherein the water storage pond is maintained at a constant year-round temperature.

* * * * *